EnGarde
Secure Professional ™

USER MANUAL

EnGarde Secure Professional 1.5

Guardian DIGITAL ™

Pioneering. Open Source. Security.

Written by Nicholas DeClario
Edited by Dave Wreski

With contributions from Ryan W. Maple, Pete O'Hara and Benjamin Thomas

Written using LaTeX

User Manual v107GD-0403

EnGarde Secure Professional User Manual

Copyright ©2000 - 2003 Guardian Digital, Inc.

Contents

1 INTRODUCTION

WELCOME TO
ENGARDE SECURE PROFESSIONAL

Guardian Digital EnGarde Secure Professional Linux is a comprehensive software solution that provides all the tools necessary to build a complete online presence, including DNS, Web, and e-mail services. EnGarde Secure Professional significantly reduces support costs due to its simplicity of use and robust security features.

EnGarde Secure Professional is a standards-based solution rich in security and Internet commerce features. EnGarde provides a comprehensive suite of applications necessary to create thousands of virtual Web sites, manage e-mail and DNS for an entire organization, manage SSL certificates, and connect high speed Cable connection, all using the integrated SSL secure Web-based administration capabilities.

This manual also includes documentation for the EnGarde Workgroup Suite, an accompanying product that was designed to provide file and print sharing capabilities, virtual private networking for remote office workers, WebMail, file and user quota abilities, as well as Windows Domain Controller support.

The Guardian Digital WebTool provides EnGarde administrators with the most sophisticated Open Source Web-based management system available. It offers secure graphical report and administration capabilities, providing the complete ability to create hundreds of virtual Web sites quickly and easily, as well as associated e-mail and DNS domain information.

1.1   Features

The EnGarde Secure Professional integrated software solution offers unsurpassed levels of security, ease of use, intrusion detection and alert capabilities, integrated database and software development packages, and support for standards-based Internet services.

EnGarde Secure Professional is also available in pre-configured turnkey rack-mount Internet servers from Guardian Digital. The Guardian Digital Linux Lock-box is a highly reliable complete eBusiness solution, configured to address space-saving considerations at co-location facilities, ISPs, and ASPs.

Guardian Digital's EnGarde Secure Professional features:

- **Browser-Based Administration** - Browser-based secure remote administration can be performed using the Guardian Digital WebTool. The GD WebTool provides security through a 1024-bit SSL connection and allows an administrator to perform 100% of the functions that could previously only be performed from the command line.

- **Guardian Digital Secure Network** - The integrated Guardian Digital Secure Network allows organizations to manage the software configuration of their EnGarde Secure Professional installations within their enterprise.

- **Web Services** - All Web functions are controllable through the GD WebTool. The creation of thousands of virtual Web sites can be easily managed and maintained.

- **High-speed Internet Connectivity** - Connect your office Cable or DSL high-speed Internet connection to build an inexpensive corporate presence.

- **Gateway Firewall Services** - The integrated gateway firewall includes the ability to protect organizations from malicious cybervandals and provides a level of assurance that its assets are secure. The port forwarding functionality provides small organizations with the ability to redirect Internet service requests to servers within the internal network. Network Address Translation provides security by masquerading requests by internal clients for Internet services as well as enabling organizations to use a single IP address for all their internal workstations to reach the Internet.

- **Intrusion Detection and Prevention** - The intrusion detection features will detect and notify you of possible threats and security related events.

- **System Logging and Auditing** - Extensive logging is performed to insure that you have the latest system information.

- **Host Security** - Security of the host itself has been significantly increased. Enforcement of longer user passwords, control of expiration dates, and utilization of the latest in advanced forms of password encryption close one of the most common and easily exploitable means of intrusion.

- **Electronic Mail Server** - The included e-mail server has been engineered to provide security and stability and can control e-mail for hundreds of domains with the click of a mouse. Mail can then be retrieved in a secure format using conventional mail clients. Additional security improvements have been made including protection from common threats as well as restricting unsolicited e-mail.

- **PHP Embedded Scripting** - The PHP HTML embedded scripting language makes it easy for developers to create dynamically-generated Web pages. PHP also offers built-in database integration for database management systems, providing the ability the produce database-enabled Web pages with a short learning curve.

- **Database Support** - The included database server provides a true multi-user, multi-threaded SQL (Structured Query Language) database server, enabling EnGarde system users and applications to create robust interactive Web sites and powerful E-Commerce sites.

- **Secured IMAP and POP3** - SSL Secured IMAP and POP3 are fully supported to help increase the security of personal e-mail.

- **Domain Name Services** - EnGarde Secure Professional can manage DNS for thousands of domains for external users trying to access virtual Web sites running on EnGarde, as well as DNS for internal users. This is all configurable using the WebTool.

- **Common Gateway Interface (CGI) Support** - The administrator has the ability to enable CGI-based dynamic Web content on an individual virtual server basis.

- **Server Side Includes** - EnGarde has the full ability to correctly display server-parsed Web pages (.shtml files).

- **Secure Shell Accounts** -The Secure Shell provides a secure encrypted communications link with EnGarde Secure Professional from a remote location, eliminating the risk previously found in other remote access methods.

- **Web Server Aliasing** - EnGarde has the ability to create thousands of virtual Web sites from the same IP address.

- **E-Mail Server Aliasing** - EnGarde gives the administrator the ability to add e-mail server aliases, allowing the creation of thousands of virtual e-mail domains.

1.2   List of Chapters and Appendices

Chapter 1 *Introduction* covers basic information about EnGarde.

Chapter 2 *General Security* gives you an understanding of basic security.

Chapter 3 *Installing EnGarde* is an guide for installing and initially configuring EnGarde.

Chapter 4 *The Guardian Digital WebTool* covers all the functions of the GD WebTool configuration utility.

Chapter 5 *Guardian Digital Secure Network* shows you how to take advantage of the Guardian Digital Secure Network automated update system.

Chapter 6 *EnGarde Connectivity* has information of the different ways of connecting to your EnGarde system from a remote location without using the Guardian Digital WebTool.

Chapter 7 The *Virtual Private Networking* (VPN) section covers configuring your EnGarde Secure Professional server for VPN and configuring Windows 98/NT/2000 to connect to a VPN using EnGarde.

Chapter 8 *Secure E-Mail* shows you how to configure different e-mail clients to work with secure e-mail services.

Chapter 9 *The Linux Intrusion Detection System (LIDS)* is covered in the WebTool but delves into a much more technical aspect of this feature.

Appendix A Quick Start Guide contains a step-by-step guide on setting up all the major components of your EnGarde system.

Appendix B The *Advanced Installer Issues* covers other features of the installer to be used by advanced users.

Appendix C *General Linux* has some basic BASH commands for getting around the system from the console.

Appendix D *Firewalls and Proxy Servers* covers how to allow your EnGarde system to get through a firewall or proxy server and how to get a client system to EnGarde from behind a firewall or proxy server.

Appendix E *Certificates* has basic information on what certificates are, how to manage them and getting a certificate signed.

Appendix F *Licenses* covers all the major licenses attached to the different software programs included with EnGarde.

Appendix G *Glossary* covers common jargon and terms used in this manual.

Appendix H *References* has a list of references used to aid in the creation of this manual.

1.3   Product Activation

Activating your copy of EnGarde Secure Professional gives you the ability to join our mailing list, priority access to the latest system and security updates and Guardian Digital technical support as described in the next section.

Activate Your Software

Guardian Digital offers the ability to activate EnGarde Secure Professional from your local desktop. Simply connect to:

```
https://www.guardiandigital.com/register
```

You can fill out all the necessary information here and submit it directly to Guardian Digital. You will have immediate access to the latest updates upon registration.

1.4 Obtaining Technical Support

Guardian Digital provides 60 days of Web, phone or e-mail support beginning at the time of product registration. This includes up to four incidents of installation and configuration support within that 60 day period. Additional support is available from your Guardian Digital sales representative.

Before contacting Guardian Digital's technical support team please visit the Guardian Digital Support Web site which covers many common technical support issues at:

```
http://support.GuardianDigital.com
```

You can contact Guardian Digital directly using one of the following means:

Phone: **1-866-GDLINUX**

201-934-9230

E-Mail: **support@guardiandigital.com**

Before you can obtain support, you must have previously registered on our Web site:

```
https://www.GuardianDigital.com/register
```

Additional details on available support plans are available at:

```
http://www.GuardianDigital.com/support
```

2 GENERAL SECURITY

Before you start using EnGarde Secure Professional we recommend you read this section covering general security knowledge. This section will help you understand the goals of your EnGarde system and in turn will help you configure it better for your needs with security in mind and increase the overall security of your network.

2.1   Why Do We Need Security?

In the ever changing world of global data communications, inexpensive Internet connections, and fast-paced software development, security is becoming more and more of an issue. Security is now a basic requirement because global computing is inherently insecure. As your data goes from point A to point B on the Internet, for example, it may pass through several other points along the way, giving other users the opportunity to intercept, and even alter it. It does nothing to protect your data center, other servers in your network, or a malicious user with physical access to your EnGarde system.

2.2   How Secure is Secure?

Security is about defense in depth. Providing physical security as well as a well-designed network, control over the users and processes on the host itself, and regular maintenance can go a long way towards providing good security.

In the most basic sense, a system is secure if it does what it's supposed to do, even if its users attempt to do something they're not supposed to do. It protects the information stored in it from being modified either maliciously or accidentally or read or modified by unauthorized users.

Consider the security of your household. Perhaps you have an alarm system, but does it work if the intruder cuts the system power? Security involves tradeoffs. How much is your data worth? Does it make sense to protect your system with the level of security you might find protecting Fort Knox, or would that cost more than the data itself? Guardian Digital provides an extremely functional e-commerce server, while still retaining all the reliability, configurability, and scalability you have come to expect with the Linux operating system.

2.3   Security Planning and Policy

Assessing risk and making prudent decisions before the system is installed is the best approach. You can go a long way towards providing good security by establishing a security policy. A security policy is a written document that outlines what is permitted behavior on the system. Once written, it is reviewed periodically and distributed to all users of the system. No system can be fully secure, but with due diligence and attention to detail, many security threats can be mitigated.

Linux is not susceptible to viruses in the strictest sense of the word (no pun intended), but permitting content to enter the system that has not explicitly been authorized will surely lead to problems.

Guardian Digital's EnGarde Secure Professional has been engineered with the greatest degree of security available on any Linux Open Source e-business server to date. No longer is it the case that a company can purchase or contract an e-commerce solution without great concern for the assurance and integrity for the data and information contained within it. Guardian Digital solutions have been engineered with security as a primary concern, providing that high degree of assurance required to conduct business on the Web today.

This high level of security integrated in to EnGarde Secure Professional requires you follow the guidelines in this manual when configuring and administering EnGarde. By following these guidelines you can be assured the highest level of system security at all times.

3 INSTALLING ENGARDE SECURE PROFESSIONAL

EnGarde Secure Professional comes with an easy to use front-end for installing the operating system. Described in the following sections are the steps to be completed to successfully complete an installation of EnGarde Secure Professional.

EnGarde Secure Professional also provides an easy to use interface for the initial configuration. The initial configuration is ran after installation to configure the software on the machine, as opposed to the installation which configures hardware. This interface requires you to configure it from another PC. The client PC can be any operating system and only requires a browser that supports SSL. Netscape 4+ and Internet Explorer 5+ will be fine for doing this.

The interface you will be using will guide you step-by-step through the set up process. We will also outline the steps in more detail in this manual. The Guardian Digital WebTool will provide the complete ability to configure your EnGarde system after installation.

3.1   System Requirements

Below are a list of the system requirements for EnGarde Secure Professional.

- 486 or faster processor

- 16MB of ram or greater

- 520MB hard drive (SCSI or IDE)

- 1 network interface card

The above listed requirements are the bare minimum for EnGarde Linux to function properly. We highly recommend using a system with the following specifications:

- Pentium class processor

- 32MB of ram or greater

- 2Gb hard disk (SCSI or IDE)

- 1 PCI network interface card

3.2   The EnGarde Secure Professional Installer

The installation process is mostly automated but can be very interactive if the advanced user wishes.

The installation process is started by booting the system with the EnGarde Secure Professional CD-ROM. If your system does not support the CD-ROM drive as a boot device you can create a bootable floppy disk, refer to *Appendix B.1 on page 272* for information on creating a bootable floppy disk.

Booting

Once the system finds bootable media you will be presented with a prompt and a few options. You can press *Enter* to continue with a normal installation, press F2 to view more information concerning *Rescue Mode* (explained in *Appendix B.1 on page 272*) or press F3 to view additional information concerning EnGarde Secure Professional and the installation process.

The Installer

Following the boot menu the kernel will be loaded and booted. Once this process is complete the installer will be launched and you will be presented with the following screen:

Here you are given the option to choose your language. Currently the installer itself does not support any languages except for English but it will accept keymappings for the languages listed. Additionally after installation language settings, keymappings and font settings will be active at the systems console.

To select your language scroll through the list with the arrow keys on your keyboard and when you your language is highlighted press 'enter' to select it.

You will then be brought to a welcome screen:

Press the *Ok* button to continue on your way.

Mounting the CD

Next you will be prompted to insert the CD-ROM. If the CD-ROM is already in the drive just hit *Ok* to continue, otherwise insert the CD-ROM and press *Ok*. It is

not necessary to close the door, it will close itself when you press *Ok*.

NOTE: Although the CD will boot from a SCSI CD-ROM drive, if configured to do
so, it will not install from a SCSI CD-ROM drive. An ATAPI CD-ROM drive
is required for installation.

3.2.1   Partitioning

The next portion of the installation process is to partition the system's hard drive(s).
The EnGarde Installer provides two methods of partitioning, *Automatic* and *Manual* methods.

For difficulty understanding any of the terms used in this section, please see the
Glossary located on page 295.

Automatic partitioning will completely partition your system for you with minimal user interaction. Manual partitioning allows you complete control over the
partitions on the system. Both modes are outlined in detail below.

Automatic Partitioning

Automatic partitioning will create the necessary partitions for you and create a
filesystem on each partition, as well as a swap partition. For detailed information
on how the drive is partitioned refer to *Appendix B.3* on page 273 .

Drive Type

The only input required from the user in automatic partition mode is to choose if
you wish to install EnGarde Secure Professional on a SCSI disk or on an IDE disk
one.

If the installer is unable to load support for your SCSI controller, you will be presented with a list from which to choose.

Note: If *SCSI Disk* is chosen your SCSI adapter must have a boot prom otherwise
             EnGarde Secure Professional will fail to boot after install.

Drive Partition Warning

Once you have made your selection a warning box will appear informing you that all data on the drive will be lost. EnGarde will install on the first drive found on the specified bus you selected. For example, if you chose IDE then `/dev/hda` will be used, and if you chose SCSI then `/dev/sda` will be used.

EnGarde Secure Professional is a server operating system and is designed to be the only system on the machine. For this reason all information on the primary installation disk will be destroyed; other drives in the system will remain untouched.

If you wish to use additional disks in your system, you must change to manual mode for partitioning.

After *OK* is selected partitioning will proceed.

Manual Partitioning

Manual partitioning mode allows advanced users to use multiple drives, both IDE and SCSI, and configure them however you like. If you don't have a clear understanding of partitioning it is recommended you use the automatic partitioning mode.

Main Screen

When you first start the manual partition mode you will see the screen on the following page.

This main screen will show you a list of created partitions, drives with space available, and space remaining. It will also allow you to add, delete and edit partitions.

There are two listboxes on this screen, the partition listbox and the hard drive listbox. Both boxes scroll and can be accessed by hitting the *tab* key on the keyboard. Hitting *Enter* while in the partition listbox will bring up an edit screen, described later in this section. To scroll up and down in a listbox simply use the arrow keys on the keyboard.

The first thing that must be done at this menu is to add a /boot partition. A /boot partition is required for compatibility and security and will be created for you automatically the first time you hit the *Add* button.

The /boot partition will be a 30MB partition created on the first drive in the system. If you have both SCSI and IDE drives in the system, the following window will appear so that you may select if you want this /boot partition on the first IDE disk or on the first SCSI disk.

After clicking the *Add* button the main screen will refresh and you will see the newly added partition. If you had to choose between SCSI and IDE this will happen after your decision.

Adding a Non-Software RAID Partition

Before EnGarde Secure Professional can be installed a / and /boot partition are required. As described above the first time you click *Add* a /boot partition is created. After that, you have the ability to create your own partitions as necessary.

NOTE: The installer will not continue until a / partition has been created.

Step 1 - Selecting a Drive

The first thing the installer requires when adding a partition is to select which drive you want the partition to be created on. It will display the following dialog showing each drive and the remaining space on that drive.

If a drive has all of its space allocated to other partitions it will not be displayed. This dialog will also not appear if you only have one drive in the system.

Step 2 - Partition Size

After selecting the drive to create the partition on you must select the size of this partition. The interface accepts input in the form of MB so for a 500 MB partition you would type *500*.

After entering in the partition size you have a second option, *Test disk integrity*. This will scan the drive for physical damage before using it. If it finds a bad portion of the disk it will ignore this portion when writing the filesystem.

NOTE: Running the disk integrity test can be very time consuming depending on the
              disk size.

Step 3 - Mount Point

The last step of creating a partition is defining where the partition will be mounted
on the system. You will need to type in the full path of the partition. You can
also choose to make this partition a swap partition by selecting the swap partition
check box.

NOTE: If you choose to make the partition a swap partition anything typed into the
              entry box will be disregarded.

Step 4 - Completion of a Partition

After selecting the mount point you will be returned to the main screen. You will
see the partition you just created in the partition listbox. Once a / partition has
been created you can:

- Continue with the installation

- Add more partitions

- Delete the partition

- Edit the partition

To delete a partition move to the partition listbox by using the *tab* key. Highlight the partition you wish to delete by using the arrow keys on the keyboard. Then using the *tab* key, select the *Delete* button to delete the partition. The partition will be removed from the listbox and its space will be allocated back to the appropriate drive.

Creating a Software RAID Partition

EnGarde Secure Linux allows the creation of Software RAID partitions. A Redundant Array of Inexpensive Disks (RAID) allows redundancy and performance over multiple hard disks. RAID is usually done by a physical hardware controller or controlled by software. If a hardware RAID controller is found in the EnGarde system Software RAID will not be an available option at installation time.

RAID has multiple configurations referred to as levels. EnGarde supports RAID levels 1 and 5.

RAID 1 A RAID 1 array consists of two hard disks and no limit on spares. This RAID level is sometimes referred to as "mirroring". It makes a mirror image of the first drive on the remaining drives. If the first drive fails a backup is used. The size of a RAID 1 partition is limited to the size of the smallest partition in the array.

RAID 5 A RAID 5 array consists of at least three disks and no limit on spares. RAID 1 offers larger partition sizes than RAID 5 with increased read performance but slightly reduced write performance over RAID 1. RAID 5 stores parity information across all disks for redundancy making it possible to recover from a failed disk. The size of a RAID 5 partition is determined by taking the total number of disks in the array minus one and multiplying that by the size of the smallest partition in the array.

RAID Spares In the event of disk failure the Software RAID system will reconstruct the RAID array using the parity information contained on its RAID disks. It will write the reconstructed data to one of the spare disks in the system. The spare disk remains unused until an error occurs. This method is sometimes refered to as "hot reconstruction". A RAID array can be fully reconstructed and operational with no system downtime.

If at least two disks are found in the EnGarde system a prompt to choose *RAID 1* or *No RAID* will be given. Additionally, if three or more disks are present in the system an option for RAID 5 will be listed as well.

NOTE: Once a Software RAID partition is created the entire system will be configured for Software RAID. Non-RAID partitions can not be created at that point.

Choosing the RAID disks and Spare disks

Once *RAID 1* or *RAID 5* is chosen a new menu with two lists will be displayed. Each list shows the hard drives located in the system and their available free space.

Two drives must be chosen from the RAID list as the main RAID partitions for a RAID 1 array and at least three drives if this is a RAID 5 array. There is no limit to the number of spare disks.

Determine size of the new partition

Once the drives have been selected the installer will determine the maximum size the partition can be with the selected drive configuration.

The size of the partition is required in MB.

Determine the mount point

The last step is to select a mount point for the new RAID partition.

Type the directory name where this partition will be located into the entry box beginning with a `/`. Enter only a `/` for the main root partition.

Once the mount point is entered the main partition screen will be displayed.

A `/boot` partition will automatically be created on `/dev/md0`, the first Software RAID partition. The size will be 30MB and it will use the drive configuration that was selected for the partition that was just created. In the partition list, below the `/boot` partition, the partition that was just created will be displayed.

Creating a Swap Partition in Software RAID Mode

The EnGarde Installer allows the selection of multiple swap partitions during a Software RAID installation. These swap partitions are assigned the same priority so that the system will access all the partitions at the same time to read and write its data. This greatly increases swap performance.

After creating the first Software RAID partition the RAID selection screen will change.

In place of the *No RAID* option will be *Swap*. Select the *Swap* option to start the process of creating swap partitions.

Selecting drives to use for swap

After selecting *Swap* there will be a new menu with a list of available drives. Choose at least one drive for the swap partition. There is no limit on swap partitions.

Determine swap size

Once the drives have been selected the size of the partition(s) must be selected. The installer will determine the maximum size allowed for the swap partition(s). This is determined by the drive with the smallest space available.

After choosing the size of the swap partitions along with any others that may have been created it will be displayed in the main partition menu.

Editing a Partition

To edit a partition move to the partition listbox by using the *tab* key on the keyboard. Once in the partition listbox, highlight with the arrow keys on the keyboard the partition you want to edit and hit *Enter*. At this point the following dialog box will appear:

You will notice you can change all the configuration choices you made when creating the partition. All the same rules mentioned previously apply here.

NOTE: /boot can not be edited.

Creating Partitions and Filesystems

Once all the partitions have been defined hitting the *OK* button will continue with the installer. The installer will display a small dialog showing each partition being created:

After all the partitions are created an Ext3 journal and filesystem will be created on each partition, except for swap.

3.2.2   Package Selections

EnGarde Secure Professional offers the ability to choose what packages you wish to have installed on your EnGarde system. In this dialog you can choose which packages you wish to install.

You can choose from the following packages:

Database Packages

> Select this option to include support for building databases Use the MySQL database to build true multi-user, multi-threaded SQL databases, enabling EnGarde system users and applications to create robust interactive Web sites and powerful E-commerce sites.

DNS Packages

> EnGarde Secure Professional can manage DNS for thousands of domains for external users trying to access virtual Web and email sites running on EnGarde, as well as for internal users. This is all configurable using the WebTool.

Firewall Packages

> The integrated Gateway Firewall includes the ability to protect organizations from malicious cybervandals. The port forwarding functionality provides small organizations with the ability to publish internal servers on the Internet. Network Address Translation provides security by masquerading requests by internal clients for Internet services as well as enabling organizations to use a single IP address for all their internal workstations to reach the Internet.

Mail Packages

> The include email server has been engineered to provide security and stability and can control email for hundreds of domains with the click of a mouse. Mail can then be retrieved in a secure format using conventional

email clients. Additional security improvements have been made including protection from common email threats as well as restricting unsolicited email.

NIDS Packages

The intrusion detection features will detect and notify you of possible threats and security related events. Select this option to enable network and host intrusion detection on your EnGarde system.

Web Packages

All Web functions are configurable using the WebTool. The creation of thousands of fully-functional virtual Web sites, including CGI, PHP, and perl support, can be easily managed and maintained. Select this option to provide services for building Web sites.

This dialog box contains a list-box that has listed all the available packages. You can navigate the list-box with the arrow keys. Once an item has been highlighted press the *Enter* key. This will turn the item red. When you move the cursor the item will appear orange. That item has now been flagged. To select another item do the same thing. To deselect an item simply highlight it and press *Enter* again.

If you choose not to install any packages just the core packages will be installed.

The only way to leave this dialog box and continue with the install is by selecting the *Ok* button.

Below is an example of what the dialog box will look like with multiple packages selected. Selected are the '*Databases*' and '*Firewall*' packages while the cursor highlights the '*Mail Services*' package.

Once you have selected which packages you wish to install you can press the *Ok* button to continue and the packages will begin to install.

NOTE: If you plan on using PPPoE you must select both the '*DNS*' and '*Firewall*' packages. Additional information concerning PPPoE, DHCP and broadband usage can be found in *Sections 4.4.8* and *4.4.9*.

As each package is installed you will see a dialog box indicating which is being installed.

3.2.3   Networking

Once the EnGarde Installer has finished installing all of the selected packages the networking configuration will begin.

The network configuration process will allow you to configure multiple network cards with static IP, DHCP and/or PPPoE configurations and set up host and domain names and your DNS configuration.

Following this dialog box the EnGarde Installer will attempt to auto detect all the network cards in the system. If any network cards fail to initialize properly a dialog box will appear. This is discussed in detail at the end of this section.

3.2.4   NIC Options

The first part of the network configuration is determine how to configure each ethernet device found in the system.

This dialog box will display in a list each ethernet device found in the system.

Following it is an '*' signifying which option has been selected. To change the option simply highlight the device and hit *Enter*; the '*' will cycle between each option.

Once you have set everything up hit *Ok* to continue. There is a *Back* button in the next dialog box if you choose to come back here and make changes.

3.2.5   NIC Static Configuration

If you selected any static devices you will be brought to the dialog to configure static interfaces.

NOTE: DHCP and PPPoE devices will be listed here but can not be edited. If you wish to change these you will need to hit the *Back* button.

In this dialog box each network card will have four categories, IP Address, Gateway, Netmask and Network. If these terms are unfamiliar to you, consult with your service provider.

IP Address An IP address is a unique number used to identify a computer on a network. Generally you can purchase a block of IP addresses you are allowed to use on the Internet, or are assigned one or more IP addresses from your service provider. Enter the IP address you want to assign to the EnGarde machine.

Gateway To give a computer the ability to talk to computers on another network they must communicate through a gateway. Enter this IP address here.

Netmask The *netmask* defines a network within the larger network, called a subnet. The netmask defines the subnet mask. Enter the appropriate subnet mask for the network, generally, `255.255.255.0`.

Network The *network* is the network portion of the IP address as determined by the network mask. For example a network mask of `255.255.255.0` and an IP address of `192.168.1.1` would denote the network address as `192.168.1.0`. This specifies the network that your server will "live" on.

When this dialog first appears default values will be inserted. Change these to your networks settings.

The following screen shot is an example after all the NIC information is entered into the system.

NOTE: If you plan on using PPPoE two interfaces are required and one must be a static IP. Additionally the static interface must be defined as the gateway and the DNS server. Additional information concerning PPPoE, DHCP and broadband usage can be found in *Sections 4.4.8* and *4.4.9*.
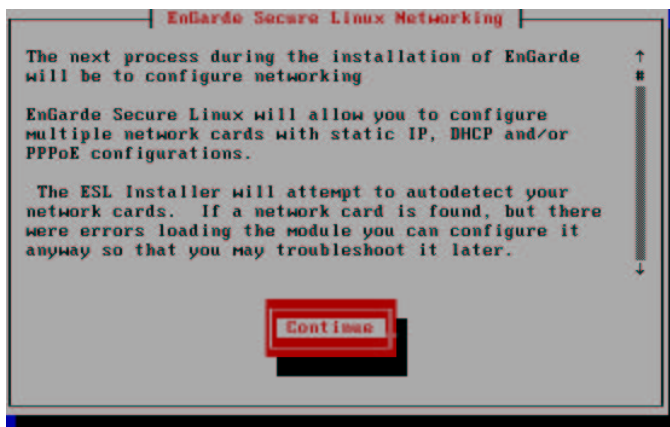
3.2.6 Set the Default Gateway

The next step in the network configuration process is to configure the default gateway. The default gateway is required; if the requested route is not found in the routing table, the default gateway will be used.

To set a device as a default gateway simply scroll with the arrow keys and make your selection by pressing *Enter*.

NOTE: When a device is configured for PPPoE it is assumed as the default gateway.

3.2.7 Configure a Fully Qualified Domain Name (FQDN)

After selecting your default gateway it's required you enter in your hostname and your domain in a Fully-Qualified Domain Name (FQDN) format.

A Fully-Qualified Domain Name is written from most specific (a host name) to least specific (a top-level domain), where each part of the domain separated by a period. For example, if you were to name the host `lockbox` and place it inside the `guardiandigital.com` domain the FQDN would be `lockbox.guardiandigital.com` as in the example screen-shot below.

3.2.8   DNS Configuration

The final step of network configuration is to configure the Domain Name Service servers. Domain Name Service (DNS) is the software that is responsible for converting host names into numbers that computers can understand.

If you selected a DHCP or PPPoE ethernet device, they may retrieve your DNS information for you. If you have no DHCP or PPPoE configured devices then you are required to enter in at least one DNS server otherwise both are optional.

3.2.9   Troubleshooting NICs

If the EnGarde Installer locates a NIC but fails to initialize the card properly you will be brought to this dialog box below at the start of the networking module.

This will list for you all the cards found in the system that failed to load properly. Since EnGarde requires at least one NIC device be present during install you will have the ability to force an ethernet device to be configured.

To forcefully add the device select the *Add* button from the dialog box. Upon doing so a new box will be displayed.

This new dialog box will allow you to choose one or more network cards to add to the system. The list-box in this dialog box allows you to select multiple network cards to be added at the once. This interface works identically to the one mentioned earlier in the package selection section.

After selecting your devices hit *Add Card(s)* and if you added every available card

you will be brought to the main network menu.

If you choose to not select all the available cards, you will be returned to the previous menu where you can choose to select the remaining cards or continue on to the network configuration at this point.

3.2.10   New User Creation

Once the network configuration is complete you are given the opportunity to add a new user during the installation process. This new user will be an administrative user, they will be part of the *'admin'* group and an SSH key will be automatically created for the user.

NOTE: The SSH key passphrase will match the users system password. This can be changed later via the GD WebTool.

If you chose to create a new user at this point, you will need to enter in the users real name, user name, and a password. Once all required fields have been entered hit *Ok* to create the account.

3.2.11   Creating a System Boot Disk

The final step in the installation process is to create a boot disk. It is highly recommended that you do so. If there are any problems with the system disks that can prevent the system from booting properly a boot disk will solve your problem.

The boot disk is to be used as a rescue tool only. It contains a kernel with minimal security installed in it so that you can fix a damaged system.

To create the disk insert a blank disk and hit *Ok*. When the disk has been created the following screen will appear:

3.2.12   Installation Complete

When this last dialog box appears the installation has completed. Remove all CD's and floppies from the systems drives and when you click *Ok* the system will reboot. Once the system is finished rebooting you can proceed to the initial configuration.

3.3   Configuring the Client Machine

A client machine is required to configure EnGarde. You will need a crossover cable to make the connection from your PC to the EnGarde machine, or you can put them both on a hub. The only drawbacks are while the system is on a hub it is vulnerable from other machines connected to that hub and the default network settings could interfere with other machines connected to that hub.

To configure you client PC you must first start by disconnecting your client PC from the network. You can simply do this by unplugging its network connection. Then change your PC's network settings. Don't forget to write down your old settings to change back to when you are finished setting up EnGarde.

Change your client PC's network settings to the following:

```
IP Address:  192.168.10.110
Subnet:      255.255.255.0
Broadcast:   192.168.10.255
Network:     192.168.10.0
```

Once you have changed your settings and the changes have taken effect, you must make sure all your proxy settings are disabled. To disable your proxy settings in both Netscape Navigator and Internet Explorer please read Appendix *D Firewalls and Proxy Servers* on page 283. Once all changes have been made to the proxy settings you will be ready to connect to EnGarde.

NOTE   Changing network settings may only be necessary if you selected the default network settings. If you configured EnGarde to work with your current network changes may both be needed.

If you have difficulty connecting after making the changes above on a Windows client, you may have to disable the *Logon to Windows NT Domain* option in your network configuration. You can do this by selecting *Networking* from the *Control Panel*, then selecting properties for *Client for Microsoft Network* and unchecking the *Logon to Windows NT Domain* check-box. You can now hit the *OK* button to finish. You may be asked to reboot your Windows system.

3.4   Connecting to EnGarde

At this point you have your client PC's network configuration set up to work with EnGarde, and you have it physically connected to your PC via a cross-over cable or both machines are connected on the same hub. You are now ready to connect to your EnGarde.

Start by powering up the EnGarde system. Next load up the browser on your PC. Either Internet Explorer 5.5+ or Netscape Navigator 4.78+ is required. First you must make certain that you have proxy servers disabled. You will not be able to successfully connect to EnGarde with proxy servers enabled. Type in the following address:

$$\texttt{https://192.168.10.100:1023}$$

It will take a few moments to connect. Once the connection is made you will be informed of a new certificate. Guardian Digital distributes EnGarde with a certificate generated by our security team. Since the certificate is not issued by a certificate authority you will be prompted to accept the certificate. Instructions on how to do this and more information concerning certificates can be found in *Appendix E Certificate* on page 289 if necessary.

After accepting the certificate you will be prompted for a login name and password. This information is pre-set to:

```
Login: admin
Password: lock&%box
```


The login and password are case sensitive. During step 2 of the initial configuration you will be prompted to change the password. You MUST change this

password. Otherwise it will remain *lock&%box*.

3.5   The Initial Configuration Process

Once you enter the login name and password you are in the EnGarde Initial Configuration.

Now we are ready to start the initial configuration of EnGarde Secure Professional. Click on the *Begin Configuration* button to start the initial configuration process.

At the main screen you will see a brief outline of the different steps you are about to be going through, each with a brief description.

From here you can start the initial system configuration. It will guide you through step-by-step. You can not skip steps here. The next section covers each step of the configuration process.

3.5.1   Password and Access Control

This first step of the initial configuration is to set the root and WebTool passwords and setup access control.

The root Password

The root password will only be used to login to the system from the console. Enter in a password that is at least six characters. Mixing numbers, letters and avoiding whole words is recommended. A few examples would be to take a word like *lockbox* and break it up with some letters and numbers. You can use the following characters as well:

!	@	#	$	%	^	&	*	(	)

So you can end up with something along the lines of:

<div align="center">

`lock%$box`

</div>

Which will be almost impossible to guess even more difficult to crack.

You have to enter the password a second time to verify they match.

The WebTool Password

The Guardian Digital WebTool password will be used every time you login to the WebTool. We suggest making this password different from the root password but still follow the password suggestions offered above.

Access Control

In this area you will have to supply a list of hosts that are allowed to access the Guardian Digital WebTool on your EnGarde system. You can list as many hosts as you want, but we recommend listing only those that are necessary for administration.

You can list them by IP address or hostname. Entering the network address will allow access to the entire network. Each item must be on it's own line.

Once you have everything filled in click *Save and Proceed* to continue with the initial configuration.

3.5.2   Locale and Time Setup

The next step of the initial configuration process is to configure the locale of your system and set up your time servers.

Locale

In the *System Locale* section you will see two pull-down menus. The first menu allows you to select your country or region. After selecting your country or region the second box will change accordingly to allow you to select a city or region found within your first selection.

Time Servers

After setting up your *System Locale* you will need to configure your NTP time servers. NTP is the Network Time Protocol and is used to keep your machines system clock in sync with the "official" time as defined by various atomic clocks.

You can select a time server from the pull-down menu or type in one of your own in the entry box.

NOTE: If you only wish to have one or two servers, please enter duplicates so that all three are filled out.

3.5.3   Firewall and Service Configuration

This is the last configuration step of the initial configuration. If you chose to install the firewall package you will have the *Firewall Configuration* section in this menu, otherwise you will have only the *Service Configuration* section.

Initial Configuration :: Email, Firewall, and Service Configuration

Email Configuration
Please enter the email address where system reports should be sent to. Several reports run nightly which the system administrator should review. These reports will be sent to this address.

Administrator Email

Firewall Configuration
Your machine only has one interface active, so host-only firewalling will be set up on this machine. If you add an (or activate) and interface at a later time, you can configure it in the **Firewall Setup** module.

Service Configuration
Below you are asked to define what services you would like active on this machine. Currently enabled services are already checked; uncheck them to disable them.

Service Name	Enable Service?
Firewalling	☑
FTP Server	☐
Web Server	☐
Domain Name Server	☐
Mail Server	☐
SIMAP Server	☐
SPOP3 Server	☐
User Password Changer	☐

Save and Proceed

NOTE: '*Domain Name Server*' and '*Firewalling*' must be enabled to allow broadband connection to work properly. Additional information concerning PPPoE, DHCP and broadband usage can be found in *Sections 4.4.8* and *4.4.9*.

Email Configuration

EnGarde Secure Linux produces nightly report summaries and other system related information. This information can be sent via e-mail to the system's administrator.

Enter the e-mail address you wish to receive these reports at in this field.

Firewall Configuration

Since the firewall package has been installed you must configure your trusted (internal) and untrusted (external) interfaces.

A list of all your interfaces will be in each of the pull-down menus.

Service Configuration

The *Service Configuration* will give you a list of all the services available on your EnGarde Secure Professional machine and the option to enable or disable them.

To enable a service click on it's check-box.

3.5.4   System Summary and Reboot

The information you entered during the Initial Configuration will now be displayed back to you for confirmation, as shown in the next screen-shot. If everything is correct click the *Reboot* button to complete the configuration process.

NOTE: Before the machine reboots you will be returned to the login screen. This is necessary for a successful system logout. You do not need to log back in.

If you used a crossover cable for configuration, remove it now and connect the EnGarde machine to your network. You are now ready to start administering your server.

4 THE GUARDIAN DIGITAL WEBTOOL

The GD WebTool is a secure on-line administration utility accessed using your browser. You have the capability to control every aspect of the system through the GD WebTool utility. In this section we will discuss the GD WebTool usage, interface, and how to take full advantage of everything it has to offer. This section does not cover using the GD WebTool for the initial machine configuration. You can find this information in the previous section, *3.5*.

NOTE: The GD WebTool is a program that is run by EnGarde. When you make changes the WebTool may take a few moments to process the changes. While this is happening your browser may report "*Host contacted. Waiting for reply...*". Do not press *back*, *stop*, or *reload* while this is happening.

4.1   Connecting and Logging into the WebTool

The GD WebTool is always running through it own personal mini Web server.
This server is securing your connection with SSL and can be accessed on port
1023. To connect to the GD WebTool program from your browser you will have
to type in the following URL:

```
https://computername.domain.com:1023/
```

We used `https` as opposed to `http`. This tells your browser you will be us-
ing an SSL secured connection to connect to the server. Where `computer-
name.domain.com` is you will replace with the actual name and domain. The
last part of the URL is `:1023/`, which specifies an explicit port rather than the
default port.

```
https://engarde.guardiandigital.com:1023/
```

This tells the browser that instead of connecting to the default port, 80 for non-
SSL and 443 for SSL connections, to instead connect to the specified port, 1023
in this situation.

If you are having difficulty connecting at this point, check the DNS settings on
your local PC or enter in the IP address instead of the hostname.

Once the connection is made you will be presented with a new certificate. Guardian
Digital issues the certificate for the GD WebTool. Since the certificate is not is-
sued by a certificate authority you will be prompted to accept the certificate. In-
structions on how to do this, and more information concerning certificates, can be
found in Appendix E *Certificates* on page 289.

Once you enter secure mode in your browser you will notice a lock that will turn
yellow. In Internet Explorer and Netscape Navigator you will see this lock dis-
played along the bottom of the browser window. Netscape will also display a
closed lock at the top of the browser. This lock will also turn yellow when in
secure SSL mode. If you click on the lock you will be provided with more infor-
mation about your current secure connection.

4.1.1   Logging in

Once the connection has been established, the GD WebTool will prompt you for a
login name and password.

Use the login name and password you specified during the initial installation and configuration of the machine. If you enter in a wrong name and/or password, return to the previous screen and you can enter it in again.

4.2   The Main WebTool Menu Screen

After a successful login the GD WebTool will bring you to the main screen:

This screen contains the main categories of options for administering your system. These categories are listed below with explanations:

Virtual Host Management This section controls Web server virtual hosts and the creation and deletion of on-line stores.

System Management System Management has all the basic Linux administration features including user control, network configuration, system time, ports and addresses settings, interface languages and SSH management.

EnGarde Auditing System The EnGarde Auditing System will give you an overview of the current running state of your system. This includes viewing user

processes, a number of different logs, current drive space, kernel information and network information.

This is quite a large section. It contains all the configuration for your Certificates, SSL connection, IP access control and the login banner.

The Guardian Digital Secure Network allows organizations to manage the software configuration of their EnGarde Secure Professional installations within their enterprise. It includes access to software updates, technical support, and security information alerts, ensuring EnGarde provides a robust platform requiring very little maintenance.

This section will allow you to create and view system backups.

4.3 Virtual Host Management

The Virtual Host Manager provides complete control over all Web server virtual host configurations. This section is also where you can create and delete an on-line store. To enter the *Virtual Host Management* section click the Virtual Host Management icon. The upper portion of this screen displays a list of virtual servers you have on your system. It has the port number, hostname and document root of that virtual host. Below that is the list of Virtual Host options.

If no virtual hosts have been set up yet, your *Virtual Servers* section will be empty. First we will discuss how to create a virtual host.

NOTE: After making any Web changes you must restart the Web server. You can restart the server by clicking the *Restart Web Server* button on the main *Virtual Host Management* page.

4.3.1 Creating a Virtual Host

In this section you will have the ability to create a *Virtual Host*, also known as a *Virtual Server*. Creating a *Virtual Host* through this method will be for hosting a Web site and will not affect any other virtual hosts. You must fill in all the required fields. A description of each field is listed below.

Address Here you can enter the IP address of your new virtual host. You are
allowed to have multiple virtual hosts on one IP address. The main reason
to do this is so you can host many sites without the need to register more
IP addresses. The Web server will know how to differentiate between the
different virtual hosts when they are requested.

Administrator E-Mail This will be the default e-mail address that will be dis-
played to a user who receives an error. Setting this to the owner and/or
system administrator of the virtual host is recommended.

Server Name This will be the name of the server. Enter in a valid FQDN.

Webmaster This is the user who will own all of the files for this Web site. You
can choose a user by clicking on "..." or you can type an existing user name
in this box.

Group This is the group that will have access to all of the files for this Web site.
You can select an existing group by clicking on "..." or you can type an
existing group name in this box.

If you wish to create a new group, click on the *Create Group* button and
create a new group. You can then select this new group using the group
chooser by clicking on "...".

Create a database for this site If this box is checked, a database will be created
for use with this site. You must enter a user name and password for access-
ing the database below.

Username If you wish to create a database for this site, this will be the username associated with accessing the database which is created.

An example username is `dbadmin`.

Password If you chose to create a database for this site, this will be the password associated with accessing the database which is created.

An example password is `gu@rd1@n`.

You can now click the *Create* button to create the virtual host.

After some processing you will be returned to the *Virtual Servers* main menu. You will see the new virtual host you created in the *Virtual Servers* list. If you created a new IP address or a new domain name for this virtual host you will have to add it to your DNS servers. Details on this are later in this section.

After the host is created you will now have the ability to edit that host.

4.3.2   Creating a Secure Virtual Host

In this section you have the ability to create a virtual host secured with SSL. Creating the secure host is similar to creating a non-secure host as was discussed in the previous section.

NOTE: If you do not have WebMail installed from the Professional Workgroup Suite the *WebMail Setup* will not appear.

Since the virtual host fields were explained in the previous section, *Creating a Virtual Host*, only the *Webmail Setup* will be discussed here.

Webmail Setup

Webmail is an interface that allows a user to read their e-mail via the web in their browser. Webmail will connect to your mail server via an IMAP connection for receiving and SMTP for sending mail. It will format messages into HTML for the user to view and respond to in their browser.

Enable Webmail Selecting *Yes* here will enable Webmail for this Web site. If this is already set to *Yes*, then by setting it to *No* you will remove the existing Webmail services, including the configuration file and profiles.

Organization Name This organization name will show up on several Webmail screens.

Domain Name This is the domain name that all outgoing e-mail will be from.

IMAP Server This is the IMAP server that the Webmail system should connect to. This should be kept as the default `localhost` unless you want to connect to an external IMAP server.

SMTP Server This is the SMTP server that all outgoing webmail will go to. This should be kept as the default `localhost` unless you want to relay email through an external mail server.

NOTE: `index.php` must be set as the document root for Webmail to work. The WebTool will set this for you.

When you are done making changes click the *Create SSL Virtual Host* button. Don't forget to create or upload your certificate for this virtual host. Instructions on doing so can be found in *Section 4.3.3 Editing a Virtual Host* on the current page found after this section.

4.3.3   Editing a Virtual Host

You can edit any virtual host settings on an existing host by clicking on the *address* of the host listed under the virtual servers.

Once you are brought to the *Virtual Server Options* page you will be presented with quite a large number of options. First, before you start making changes, check at the top of the page, below the Guardian Digital banner, you will see a list of options.

Make sure you are editing the intended host. In place of `lockbox.guardiandigital.com` will be the name of the site you are editing.

The options in this section are for advanced users who have knowledge of the Apache server. There are many complex options to give you full and complete control over your virtual host. We recommend you read the main Apache documentation, which can be found at `http://www.apache.org/docs`, before making any changes. There are also numerous books available on this subject.

Networking and Addresses

In this section you will have the ability to define what interfaces and addresses this virtual host should listen on.

First you will need to enter in the server administrators e-mail address. Following that is the *Alternate virtual server names* section. You have the ability to assign other names to your host. For example, say you have `www.guardiandigital.com` and you also want `www.guardiandigital.net` to go to `www.guardiandigital.com`. You would enter `www.guardiandigital.net` into the *Alternate virtual server names* field.

Click the *Save* button to save your changes.

Document Options

Here you have the option to configure specific Apache settings for the specified host.

Server-side includes and execs This will give you the ability to turn on server side includes and allow CGI scripts to be executed within them. Server-side includes are modules or programs that run on the server. CGI and Perl scripts are both server-side includes because they run on the server, while Java and JavaScript are executed on the client.

Server-side includes This works the same as the above option except it turns off the ability to execute CGI scripts.

Generate directory indexes With this option enabled Apache will create a file index when a directory is specified from the Web browser. It will create a clean list of files, with modification dates and file types.

Error Handling

Error handling is what the Web server does in the event a request is made resulting in an error. For example, if you try to go to a page that doesn't exist on a server you will see the all too common "*Error 404: File not found*.". In this menu you can list the error number and tell Apache to load a specified Web page or display a specified message if this error is encountered. Below are a list of common error codes and their meanings. You can refer to the Apache documentation for a complete list of error codes.

Error Code	Meaning
301	Permanent Redirect
302	Temporary Redirect
401	Bad Password
403	Forbidden / Access Denied
404	File Not Found
405	Method Not Allowed
500	Internal Server Error

Aliases and Redirects

This section allows you to set up aliases and redirects. A brief explanation of the differences between redirects and aliases is a CSR is a request for a signed certificate you can give to a Certificate Authority to sign. given to avoid confusion.

An *Alias* allows documents to be stored in the local file system other than the defined document directory. When a user accesses a document through this alias it will appear in their browser as if it was in the aliased directory, keeping the actual directory hidden from the user. This can be useful when you don't want a user to know where they really are or to have links and URL references that have a "clean" look. For example if you have files stored in:

```
/home/httpd/html/updates/december/2000/documentation
```

you can alias the address to:

```
/home/httpd/html/documentation
```

allowing you to keep everything organized neatly on your server while keeping the URL short for the user.

For the example given above you would need to type in:

```
updates/products/december/2000/documentation
```

in the *From* field and type in

```
documentation
```

in the *To* field.

NOTE: When setting up an alias the path is relative to the document path setup in the Web server.

A *Redirect* maps an old URL into a new one. The new URL is returned to the client which attempts to fetch it again with the new address. The browser is aware of this new address and will be visible to the user in the URL location field in their browser. This could be useful if you wish to point the user to another server. An example of this could be if you are moving a page:

```
http://www.guardiandigital.com/documentation/october
```

to another directory on your web site. In this example we are redirecting documents dated from October to the archives section of the website,

```
http://www.guardiandigital.com/doc/archives
```

Using the example given above you would need to type in:

```
documentation/october
```

in the *From* field and

```
doc/archives
```

in the *To* field.

NOTE: As with aliases above, the redirect paths are relative to the URL.

Hopefully you have a clearer understanding between the differences of aliases and redirects. In this section you will see two fields, *Document directory aliases* and *URL redirects*.

Document directory aliases This will allow you to alias a new document root. Enter the directory you want the user to see in the *From* field and where it will actually be pointing to in the *To* field.

URL redirects This will allow you to map one URL on to another. Simply enter in the original URL and where you would like it to point to. The source and destination must both point to valid URLs.

Directory Indexing

This section defines the initial page when the Web browser client requests a URL without specifying an explicit filename. For example, if you type in `www.guardi andigital.com`, it is really loading `www.guardiandigital.com/inde x.html`. If the Web server doesn't find an index file it will return a directory listing. Generally `index.html` or `index.htm` is used. You can specify more than one.

Certificate Management

There are two types of certificates: "self-signed" certificates and "signed" certificates. A "signed" certificate is issued by a Certificate Authority (CA) such as Verisign or Thawte. A "self-signed" certificate is simply a certificate that has not been issued by a CA. This provides the authentication part of the process, because the certificate has been signed by an external authority.

All of the certificate management can be done in the WebTool. You should not do any of this by hand unless you have a very good idea of what you're doing, since if it is done incorrectly it will cause the Web server to fail. As was said above, the certificate and key are a pair. If for some reason the certificate and key that are in place do not "match" each other then the Web server will fail to start. If the Web server fails to start then all of the other sites on the machine are inaccessible.

The *Certificate Management* section will allow you to configure your SSL certificate. This option will only be available if the virtual host you are editing has SSL enabled. Once at this menu you will be presented with three options which are each discussed below.

Generate Certificate and Key

Here you will see a screen similar to the certificate generation screen when creating a virtual host. All the fields are required. Upon completion of this form you a self-signed certificate and key pair will be created for the site. A description of each field is given below:

Authority Name The authority name is the name the server the certificate will be used on. For example `www.guardiandigital.com` or as in the example above, `lockbox.guardiandigital.com`.

E-Mail Address The e-mail address for the contact in control of this certificate should be entered here. An example would be `ca@guardiandigital.com` or as in the example above, `admin@lockbox.guardiandigital.com`.

Department Here you can enter in the name of the department this certificate will be used in. An example would be *E-Commerce*.

Organization This is the name of the organization who owns the certificate. In the example above *Guardian Digital, Inc.* is used.

City This field requires you enter the name of the city in which the organization resides. You must enter in the full name of the city. In the example above *Upper Saddle River* used.

State or Providence Here you will need to enter in the state in which your organization resides. You must enter the full name of the state, not an abbreviation. In the example above *New Jersey* used.

Country Enter in the country in which the organization resides in this field. This
requires an abbreviated name for the country, not the full name as in the
previous two fields. In the example above *US* was used.

When all the fields are completed click the *Generate Key* button to create the
certificate and key. You must now go back to the previous screen and click the
Restart Web Server button for the changes to be activated.

Generate Certificate Signing Request

A Certificate Signing Request (CSR) is what is sent to a Certificate Authority
(CA), such as Verisign or Thawte to request a signed certificate for your site. This
section will allow you to create one to be submitted. The form looks similar to the
Generate Certificate and Key form above. You can refer to the previous section
above, *Generate Certificate and Key* for a description of each of the fields.

There is however, one new field, *Create New Certificate/Key Pair*. If this option is
selected it will create a new certificate and key with the information you filled in.
It will then allow you to download the certificate to be signed. If you wish request
a new certificate because your old one has expired then d not select the *Create
New Certificate/Key Pair*.

NOTE: This new certificate will not be used on the site until you upload it. It is meant
to be signed by a Certificate Authority.

Certificate Generation
For lockbox.guardiandigital.com
This form is to create a new certificate signing request (CSR) for lockbox.guardiandigital.com. A CSR is used to pass along to a certificate authority (CA) such as Verisign or Thawte to produce a signed certificate for this site. For more information on getting a certificate signed please refer to the documentation.

If you do not have an existing certificate/key pair, check the "Create New Certificate/Key Pair" box. If you do have an existing certificate/key pair then the CSR will be generated using the existing key. When you get the signed certificate back from the CA you can simply drop it into place for use with the existing key.

Please note that by checking the "Create New Certificate/Key Pair" box, you will overwrite the existing certificate/key pair, if any.

Create New Certificate and Key
☐ **Create New Certificate/Key Pair**

Certificate Signing Request

Authority Name	www.guardiandigital.com
E-Mail Address	admin@guardiandigital.com
Organization	Guardian Digital, Inc.
Department	
City	Upper Saddle River
State or Province	New Jersey
Country	US

Generate CSR

Once you have all the fields filled in you can click the *Generate Certificate* button and you will be presented with your certificate.

Enter Certificate and Key

If you already have a certificate and a key or have sent a CSR to a CA and have received the signed certificate back, then you would want to upload it here from your local machine. This section will present you with your current SSL Certificate and give you the ability to upload a new certificate and key.

If you have a certificate and key in place then it shows you four things:

Fingerprint: This is the unique ID of the certificate

Valid: This is the data range for which the certificate is valid.

Subject: This is who the certificate is fore

Issuer: This is who has signed the certificate.

Clicking the *Browse...* button will allow you to browse through the files on your local machine and select the certificate and key. You can then click the *Save* button to save the certificate and key to the server.

WebMail Configuration

If you chose to add Webmail capabilities to this virtual host, then the following screen will be active to allow you to make changes.

In the first section, Webmail Status, the Guardian Digital WebTool will tell you the current running status of Webmail and the URL to access it.

Following that is the Webmail Configuration which has all the options presented to you in the initial creation of the virtual host. All the options are described previously in Section 4.3.2 on page 58.

Server Configuration

Here you can alter the basic virtual host settings. You have the ability to change the IP address of your virtual host and the server name of the virtual host. You can also delete the virtual host and change the database password from here.

4.3.4   Web Site Directory Structure

When a Web site is created, the following directory structure will be created on the system:

```
/home/httpd/<sitename>-<port>
```

Inside of this directory, the following sub-directories will exist:

cgi-bin This is the directory where /cgi-bin/ is aliased to.

html This is the document root.

logs This is where the access, error, and SSL logs are kept.

If a secure site was created, the following will also be created:

ssl This is where the SSL certificate and key are kept.

cgi-bin The CGI files for you Web site should be located here. For ex-
ample, if register.cgi was placed, then you would access it by
using the following URL:

```
http://www.engardelinux.com/cgi-bin/register.cgi
```

Using the `lockbox.guardiandigital.com` example being used in this section the directory URLs would look as follows:

For a standard, non-secure Web server:

- /home/httpd/engarde.guardiandigital.com-80/cgi-bin

- /home/httpd/engarde.guardiandigital.com-80/html

- /home/httpd/engarde.guardiandigital.com-80/logs

- /home/httpd/engarde.guardiandigital.com-80/ssl

For a Secure Socket Layer (SSL) Web server:

- /home/httpd/engarde.guardiandigital.com-443/cgi-bin

- /home/httpd/engarde.guardiandigital.com-443/html

- /home/httpd/engarde.guardiandigital.com-443/logs

- /home/httpd/engarde.guardiandigital.com-443/ssl

In an HTML form, you would use something of the sort:

```
<FORM ACTION="/cgi-bin/register.cgi" METHOD="GET">
```

html

This is where the HTML files are kept.

logs   This is the directory where the logs are kept. You can set up how often the logs are analyzed in the *Configure Website Log Analysis* section of the WebTool.

ssl

If this is a secure site, then this is where the certificate and key are kept. You should never edit anything in this directory by hand.

4.3.5   Set Up Name Virtual Hosts

A Virtual Host has to be bound to an IP address. This is required for proper operation of your virtual host.

Here is where you can enter in the IP address and port of your new *Name Virtual Hosts*.

To add a new host select the port from the pull-down menu and enter in the IP address you want. The port pull-down menu gives you two selections. Port 80 for normal connections and 443 for secure connections. Choose accordingly. Click the *Add New IP* button after each IP address your your new host will be added.

To delete a named virtual host simply click on the IP address of it.

4.3.6   Configure Web Site Log Analysis

Each virtual host running on your system has it's own status logs. In here you have the options to configure these logs. You will first be presented with a list of the existing non-SSL virtual hosts. Select whether you would like to have the Web statistics generated daily or weekly.

In this menu you will have the following options:

Site Name Here you can enter in the name you wish to associate with this site. Leaving it as the name of the virtual host is a good idea.

Frequency The Web statistics software can be run daily or weekly. It's up to you how often you want new statistics generated.

Click the *Save Settings* button when you've finished your selection.

Going to the site name followed by WEBSTATS will display the logs for your virtual host. Using the example above, you would type in:

```
http://engarde.guardiandigital.com/WEBSTATS
```

User Access Control

Web statistics are protected so no one can view them without a user name and password. Since, most likely, your Web statistics are private information you will want to protect the Web statistics from unauthorized visitors. Here we will assign user access control.

Here you have two fields, *Username* and *Password*. This allows you to assign a username and password to your statistics directory. When a person tries to access them, a username/password window will appear. This allows you to define who is authorized to access your log statistics.

NOTE: By default no users have access.

4.4   System Management

The *System Management* section contains all the system configuration options for
administering the system. On the main screen you are presented with a list of all
the user accounts.

Following this section is the *Service Configuration* section and then the *System
Configuration* section.

We will discuss the user accounts portion first..

4.4.1   User Account Administration

In this section we will describe how to add users, delete users, edit users, and configure groups. These are the regular system users. Users who wish to have SSH access to the machine will need an account here. For more information on users and groups refer to the *Groups and Users* section in *Appendix* C.5 on page 281. You should see all users listed in the table, as follows:

Create a New User

To create a new system user start by clicking on the *Create New User* button. This will bring you to this screen:

Here you will enter all basic user information. Below is a brief description of each option:

Username Enter a unique user name in here. A username can not contain spaces

or special characters and can be no more than 16 characters in length. For example:

User name	Valid	Reason
Nick DeClario	No	Contains spaces
nick	Yes	<16 characters and no spaces
Nicholas DeClario	No	>16 characters and spaces

Real name The users real name. This will be the real name of the user. You can enter in their full name. Using the example above, *Nick DeClario* would be valid.

Password Enter in a password for the user. This password will be asked if the user logs into the console or needs to retrieve their e-mail.

Access Enabling this will allow a user to only access their e-mail via a secure IMAP or POP3 client. This will prevent the user from physically logging into the machine.

Windows Password Entering a password in this field will grant the user Windows File Sharing access. This password will be used for logging in to shares and domains.

Now we must set up the user in a group. Read the *Groups and Users* section in *Appendix* C.5 on page 281 for more information on user groups.

Primary Group You either can create a new group for this user or use an existing group.

Secondary Group If you want this user to additionally be part of another group you can choose that group here.

We are now ready to create the user. Press the *Create* button. You will be brought back to the main *System Management* page indicating the user has been created successfully.

NOTE: When creating a new user that user is automatically given their own private group. For example, user *nick* will automatically be given group *nick*. This allows user *nick* to have private files that no other user but root can access.

Edit a User

To start editing an existing user, click on the user name from the main *System Manager* menu. You will be brought to the same screen as for creating a new user, except it will contain all the information about the user you selected. From here just change what you wish to change and select *Save* when done. The options work exactly the same as creating a new user in the previous section.

Configure Secondary Group

The last user option in this section is the *Configure Groups* option. In here you can create and change the group names, group ID's, and members by selecting the *Configure Groups* link to edit the groups.

You will be presented with a menu listing all the current groups and giving you an option to create a new group.

Create a New Group

Selecting the *Create a New Group* link you will be brought to a new menu to create a new group.

The interface will assign a group ID. It is advisable that you leave the default value. You will also need to assign a group name and select users to this group, if necessary.

Once all the fields have been filled out hit the *Create* button to apply the new changes.

Edit an Existing Group

Editing an existing group allows you to change the group ID and what members are part of the group.

If you change the group ID you will see three options at the bottom of the menu concerning changing the group ID on files. If you changed the group ID and select no then files belonging to that group will still contain the old group ID.

Selecting the *Home Directories* option will change only files in users home directories while *All Files* modifies every file on the system in that group.

To delete the selected group click the *Delete* button.

The reason to change a users group would be to change their privileges. For example, if you want a certain user to be able to administer your EnGarde system you may add that user to the admin group. Perhaps you want a certain user to only be able to edit their own personal files and the Web files, you may add them to the www group. A brief explanation of the groups in the example above is explained below:

admin   The *admin* group will give a user access to some of the systems services. This would be good if you have other trusted users whom you wish to do administrative tasks such as maintenance, file cleanup and other needed tasks.

users   This is the group general users would be put in for e-mail access and basic system access.

mysql   The mysql group is primarily used for running the MySQL server. This is done for the same reasons as explained above in the admin

description. The administrator will also have access to MySQL and
all its databases.

named   The named group is used for the DNS server. This group is defined
specifically for this task. By giving the DNS server it's own group
helps increase security.

snort   Snort exists for the same reasons the *named* group exists.

4.4.2 FTP Configuration

EnGarde Secure Professional includes a secure FTP server. You can configure
your FTP server from here. *Global Configuration* makes system-wide changes
and the *Define Chroot* and *Blacklist* menus allow you to define who is not allowed
to connect via FTP and where users are limited to.

Global Configuration

The Global Configuration section allows you to make system-wide configuration
changes to your FTP server. Each item found on this menu is explained in detail
below.

Allow Anonymous Logins Enabling this feature will allow anonymous user logins. All anonymous users will be chroot'ed[1] to /home/ftpsecure.

Allow Local Logins This will allow local users to FTP into the machine, assuming they are not on the blacklist. A local user is defined as being a user that has an account on the EnGarde machine.

Chroot All Local Users This will chroot all local users to their home directory. When a local user logs into via FTP they will be placed in their home directory.

Enable User Uploads Enabling this will allow local users to upload files. By default local users can only download files.

---

[1]chroot is a program that will put the user in a pseudo filesystem, sort of like a jail. This will prevent the user from being capable of accessing the rest of the system but still have functionality.

Allow Anonymous Uploads Enabling this will allow anonymous users to upload files. It is strongly recommended you do not enable this ability

Allow Anonymous MKDIR By default anonymous users can not create directories. Enabling this will override this functionality.

Create Permissions (Users) Setting this option to *Owner Readable* will make all uploaded files accessible only to the user who uploaded them while the other option, *World Readable* makes files readable by all users.

Create Permissions (Anonymous) This works the same as setting the permissions for local users, as described above accept it applies to anonymous users.

FTP Banner This is a text banner that is displayed to the user when they login via FTP.

Interface to Listen On Select which interface you want the FTP server to accept connections from. Leaving this blank allows connections from every interface.

Max. Rate For Anonymous Users This is the maximum data transfer rate permitted, in bytes per second, for anonymous clients. Set this to 0 for unlimited

Max. Rate For Local Users This is the maximum data transfer rate permitted, in bytes per second, for local authenticated users. Set this to 0 for unlimited.

Define Chroot and Blacklist

This page allows you to define what users should be chroot'ed. Any user not listed here will not be chroot'ed unless you have enabled *Chroot All Local Users* in the *Global Configuration* section.

The blacklist defines what users are not allowed to FTP into the machine. If you have *Allow Local Logins* enabled in the *Global Configuration*, and you would like to block access to certain users, select their username here.

NOTE: Any changes made here will take effect immediately after pressing *Save Changes*.

4.4.3   Secure Shell Management

Secure Shell (SSH) is a program for logging into a remote machine and for executing commands on a remote machine. It is intended to replace `rlogin` and `rsh`, and provide secure encrypted communications between two untrusted hosts over an insecure network.

This section will allow you to edit the SSH configuration, generate a new host key and generate user keys.

Edit the SSH configuration

By clicking on the SSH Configuration icon you are brought to the *Edit SSH Configuration* page. Here you have the ability to allow and deny groups and users SSH abilities. Be careful when editing these options since you may grant access or deny access to the wrong people, which could cause problems.

By default EnGarde Linux will not allow you to login via SSH as the root user. Though if this feature is required it can be enabled by selecting '*Enabled*' from the pull-down menu.

The second field contains the option to define which interfaces SSH will listen on. Leave this field blank to allow it to listen on all interfaces or enter in each interface, by IP, using a blank space for the delimiter. You can also select the '. . .' button to bring up a list of all the interfaces.

In each deny/allow field you can enter in a group name or user name, whichever is appropriate for the field, using a blank space as a delimiter. Clicking on the '. . .' button will bring up a small window containing a list of users or groups you may select from.

There are a few rules to take note of when configuring access control for SSL. Below is a short list of basic rules:

- Once you add a user or group to the *Allow* sections, all other users that are not listed will be denied.

- If you add a user to the *Allow Users* section but the group the user belongs to is in the *Deny Groups* section, the user will be denied access.

- The deny rules take precedence over the allow rules.

- You may deny a user but allow the group the user belongs to.

Most configurations will be safe allowing the *admin* group access. This will automatically deny everyone else who is not part of the *admin* group.

After you have finished making your changes click the *Write Configuration* button for the changes to be saved.

SSH Key Management

The *Key Management* section allows you to create new SSH keys for your users.

Generate a user key

Generating a user key will allow your users to log in to your EnGarde system remotely via SSH. First click on the *Generate User Key* button. This will bring you to a new screen with a form to be filled out. It first requires a user name. You can type in the name or select it from a list by clicking the "..." button.

An IP address is not required but recommended for increased security. The IP address will tell EnGarde from where this user is authorized to connect. If you do not enter in an IP address it will let this user connect from any IP address.

The description field allows you to enter in a short description. This description will be displayed back to the user every time they attempt to connect to EnGarde using an SSH client such as MindTerm. For more information concerning MindTerm read *Section 6 EnGarde Connectivity* on page 179.

Finally you need to enter a password. Select any password that is at least 5 characters. Now click on the *Generate key* button.

You will now see a screen with the results of the SSH Key generation.

You now have the option to download your public key. You will need to have a copy of your key to load into your SSH program to so you will be able to gain access to the machine. Save the file in a secure location.

4.4.4   Mail Server Management

The Mail Server Management section will give you complete control over your mail server, giving you the ability to add/remove users and aliases and other mail options.

On the main menu you will have four main options, *Mail Server Configuration*, *Domain Management*, *Mail Routing* and *Stop Mail Server*.

Mail Server Configuration

Here you have the option to set up various system-wide options.

Send outgoing mail via host The *Deliver directly* option will forward any outgoing mail not destined for users of your system directly to the given host.

> If the mail server is behind a firewall or proxy server to the outside world, you will need to tell the mail server where to forward non-local mail. You can enter in a hostname or IP address here.

Allow Incoming Mail By default (Enabled), the mail server can both send and receive mail. If this is set to *Disabled*, the machine will not be able to receive mail (but will still be able to send mail).

Enable Procmail *procmail* is a mail preprocessor. When a message comes into the machine, it is passed to procmail which then looks for a file called ".procmailrc" in the recipients home directory. This file can contain "filters" to file the message into mailboxes.

> Procmail is enabled by default. If you would like it disabled, you can do so here.

Domain Management

The *Domain Management* section allows you to create a new mail domain, explained below, and to edit an already created domain. Creating a new domain is

quite simple. Below the *Domain Management* menu you will see the *Create New Domain* menu. Here you have two options, *Domain* and *Postmaster*. Both fields are required.

Domain The domain is simply the name of the domain you wish to receive mail for. For example, if you wish for the mail server to receive mail for guardiandigital.com then you would enter guardiandigital.com into this field.

Postmaster If a user sends an e-mail to a non-existent account it will be forwarded to this user. It's an administrative address that receives all undeliverable mail.

Creating a Domain

To make changes to a domain you have created you can simply click on the domain name listed under the *Domain Management* menu. This will present you with the following screen.

To create the virtual domain start by entering the domain name into the *Domain* field followed by the postmaster's address for this domain in the *Postmaster* field. Clicking *Add New Domain* will create this domain.

Editing a Domain

Once a domain is created you will see it listed under *Domain Management*. Clicking on the domain name itself will allow you to edit its attributes and add users to this domain.

To add a user, give the user an e-mail username in the *E-Mail Username* field and fill in the real user's e-mail address in the *Recipient* field. Click the *Add New* button to add this user's e-mail.

Additionally, towards the bottom of this menu, the current configuration can be changed from here.

Mail Routing

The mail routing section allows you to select what domains you would like aliased. If you have a user at the guardiandigital.com domain, and want every user to be able to receive mail to linuxsecurity.com as well, this menu provides that ability. Refer to *Figure 3.8.8*.

Enter in the domain you want the mail aliased as. We used linuxsecurity.com to create an existing mail route in the above image. We then enter in the *Relay mail to...* field the actual domain the mail should go to, guardiandigital.com in this example.

Figure 1: 3.8.8 - Mail Routing

NOTE: Subdomains are automatically included in the route.

Select the *Add New* button and the new options you entered in will appear in the *Existing Mail Routes*. Click the *Save* option to save or the *Delete* button to delete a mail route.

4.4.5   DNS Management

The *DNS Management* section will allow you to fully configure your EnGarde system's Domain Name System (DNS) settings. You will be able to add and delete master and slave zones and have the ability to edit all global options.

The Domain Name System (DNS) is the software that is responsible for converting hostnames into numbers that computers can understand. For example, the name www.guardiandigital.com corresponds to the host IP address 63.87.101.80 and vice versa. The *DNS server*, sometimes called a *name server*, is the process that runs on EnGarde Secure Professional awaiting incoming name service requests.

For example, if the DNS server is given an IP address of 63.87.101.80, it will look it up in a database of addresses and link it to it's domain name. In this example 63.87.101.80 will resolve to www.guardiandigital.com. DNS will also work the other way. Giving it www.guardiandigital.com will result in 63.87.101.80.

Before you can configure your own DNS server, you must first register your DNS server and domain name with Network Solutions or another naming authority by completing their host registration form. You will need to reserve one IP address for use by your nameserver. In order to maximize availability, every domain must have both a primary and secondary DNS server, and both must be registered with a naming authority such as Network Solutions. Guardian Digital can assist you with this process if you wish.

The *DNS Management* section contains three options, as shown below.

This section provides the ability to:

Global Option Forwarders and other various defaults that will apply to all the zones you manage.

Create a New Master Zone This will bring up the configuration screen to create a new DNS master zone

Create a New Slave Zone This will bring up the configuration screen to create a new DNS slave zone

Create a New Master Zone

The domain namespace is divided into regions called zones. For the purposes of this document, it is sufficient to describe a zone as a domain, or section thereof, for which the server will be responsible. The host `www.guardiandigital.com` is a member of the domain `guardiandigital.com`, as is `mail.guardiandigital.com` and `dns.guardiandigital.com`.

For example, *Figure 3.8.6a* shows the guardiandigital.com zone and two hosts within the zone.

Figure 2: 3.8.6a - Example of the guardiandigital.com zone.

When you select the option to create a new zone you will be presented with the page in *Figure 3.8.6b*.

The above page has quite a few options. Here we will discuss each one in detail.

Zone type The zone type will allow you to choose between forward and reverse lookup.

- Forward lookup allows the client machine to supply a Fully-Qualified Domain Name (FQDN) and the DNS will return the IP address.

- Reverse does the exact opposite. You supply an IP address and the DNS will return an FQDN.

NOTE: When creating entries for a *Reverse Master Zone* you must **not** put in entries that refer to an alias. To do so would *break* DNS for the corresponding domain.

Domain name / Network This contains the actual domain name, or, in the case of reverse zones, the network address block, that this DNS zone will reside in. For example, if your EnGarde system is like above,
`lockbox.guardiandigital.com,` then the domain would be
`guardiandigital.com.`

Figure 3: 3.8.6b - New Master Zone Options

Master Server This section will contain the IP address of your master DNS server. The master DNS server, also known as a *Primary DNS Server,* maintains a list of domain names and their IP addresses. This list is made available to other DNS servers on the Internet so that users can access these sites over the network. For example, if you own `guardiandigital.com` your master server will control `guardiandigital.com`. You can have other DNS servers, known as *secondary DNS servers,* or *slave DNS server*s, that act as a backup to the primary DNS server for `guardiandigital.com`. If your EnGarde system is your master DNS server then enter in the address of your EnGarde system.

Email Address The default e-mail address associated with this zone. Generally this is the e-mail address of the system administrator or whomever is responsible for DNS on your network.

Allow Transfers From... DNS will need to transfer information if you have slave DNS servers on your network. This should contain a list of IP addresses and/or a block of IP addresses for other DNS servers that are allowed to transfer DNS information between each other. You can set the default in the *Default Zone Settings* section for this specific zone, which is described later in this section.

Allow Queries From... Here you can list the IP addresses and/or block of IP addresses for machines that are allowed to query your DNS server. You may want to limit this to the people inside your network if your EnGarde system is located on your internal or private network. We recommend leaving the default set if you are uncertain. You can set the default in the *Default Zone Settings* section, which is described later in this section.

Creating a New Slave Zone

A secondary DNS server, also sometimes referred to as a slave server, for a zone gets the zone data from another DNS server that is authoritative for the zone, called its master server. When a secondary name server starts up, it contacts its master server and requests a copy of the zone data for which it is responsible. This is called a *zone transfer*.

A slave server will backup your master server. This is mostly for redundancy if your master server is not running or is unavailable to answer a query. This section has everything necessary to create one.

NOTE: You must configure the master server to allow this new slave server to per-
              form zone transfers from the master server. These changes must be made on
              the master server. Information pertaining to this can be found in *Section* 4.4.5
              Edit Master Zone on page 101.

The options on this screen are the same as setting up a master server. Find the
detailed information in the previous section.

However, there is one new category, *Master Servers*.

Master servers In the master servers section you can list all the master servers
              that this slave server will obtain its DNS information from. At least one
              master server is required in this section.

NOTE: You are required to list your slave server as a name server on your master
              server. You can find information on doing this in the *Name Server Section* on
              page 104.

To finish creating a new slave zone you will need to define a mail route to backup.
Defining a mail route must be done from the master server. You will need either
the Fully-Qualified Domain Name (FQDN) or IP address of the slave server that
will be handling the mail route. Information on configuring this on your master
server can be found on page 106.

A New DNS Management Screen

Once you have completed the zone creation form, click the *Create* button. You will be returned back to the main screen. Now you will have a list of options at the top, followed by a list of your DNS servers.

The first object in this menu is the *Global Server Options*. Here you have the ability to create new Master and Slave zones, discussed above, and to edit the *Global Options*.

Global Options

Global Forwarding and Zone Transfer Options

Servers to Forward Queries to... A forwarder is used for name servers that may not necessarily be directly-connected to the Internet. This may be due to being behind a firewall, or inside of a corporate network. Forwarders will instead query a specified additional name server for its DNS information. If your DNS server will be responding to a forwarding server you will want to specify the server(s) it is allowed to contact. See *forwarders* and *forward zone* in the glossary for more information concerning forward queries.

Addresses to listen on This allows you to define which address your want your DNS server to listen on. You can enter in each IP address by hand, leave the field blank for it to listen on all interfaces or use the '...' button to select the interfaces from a menu.

NOTE: A forward server is still a primary or slave server; don't get confused here. All outside queries will be given to it first.

Default Zone Settings

Allow transfers from... This sets the servers that are allowed to perform zone transfers from the DNS server. When a slave server requests updated information from the master server, the master server will transfer it to the slave server if authorized. This procedure is known as a *zone transfer*. No servers are authorized by default. If you are uncertain of what to enter in here, leave the default set and contact your network administrator.

Allow queries from... This sets from which IPs your DNS server will accept DNS queries. By default the DNS server will accept queries from all IP addresses. If you are uncertain about what should be entered in here, leave the default on.

Existing DNS Zones

The other section on the main DNS page below the *Global Server Options* is *Existing DNS Zones*. This will display the reverse and forward addresses of a domain. If you click on the address you will be brought to the corresponding options page to have the ability to make changes. The reverse address page and the forward address page both have different options. We will discuss both pages below.

Edit a Slave Server

In this section you have the ability to make changes and delete a slave server. You should be familiar with these options since they were used to create the slave server and in the *Global Options* section. Refer to those sections for more detailed information.

Edit a Master Zone

Add Address Record

The *Address* section will allow you to define address records. In the given address (i.e., `smtp.guardiandigital.com`) you can define specific servers. The menu is broken down into two sections, *Add Address Record* and a table of the current records listed by IP address followed by the hostname. Take note that these records are only valid for the defined zone.

To create a new Forward Address Record you simply need to fill in the two required fields described below.

Hostname The hostname is the Fully-Qualified Domain Name (FQDN) for the specified machine.

Address In the address entry field you will need to enter in the IP address of the machine for this record.

Create "Default A Record" Check this box to make this new address record the default A record.

Once you have filled in all the fields you can click on the *Create* button to create the new forward address. Once the page refreshes you will see it listed at the bottom of the page.

Edit/Delete a Record

Once a record has been created and you see it listed below the *Add Address Record* menu, you will have the ability to edit the record by clicking on the name of it. This will bring you to a new screen that is similar to the *Add Address Record* screen.

To edit the name server simply make your changes directly in the *Name Server* field and click the *Save* button to make the changes. If you wish to delete this name server record click on the *Delete* button.

Name Alias

The *Name Alias* section gives you the option to configure an alias for this record.

On this menu you have two options, *Alias* and *Real Name*.

Alias The alias needs to be a Fully-Qualified Domain Name (FQDN). In this case the alias is where you want the user to be redirected to. For example, the user types in `www.guardiandigital.com` while really they are being sent to `lockbox.guardiandigital.com`.

Real Name The real name of the server also needs to be a Fully-Qualified Domain Name. This is the name that the Alias will really be going to. In the example above you would enter in `lockbox.guardiandigital.com`.

Edit/Delete an Alias

Once you create a new alias it will appear at the bottom of the page.

Similar to the other sections, you can click on the name to edit the record. After clicking on the name you will be brought to the *Edit Name Alias Record* page.

You can make your changes by editing the appropriate field. When you are done with your changes you can click the *Save* button to set the changes. To delete the record simply click the *Delete* button and the alias will be deleted.

Name Server

The Domain Name System (DNS) is the software that is responsible for converting hostnames into numbers that computers can understand. For example,

the name `www.guardiandigital.com` corresponds to the host IP address `63.87.101.80` and vice versa. The *DNS server*, sometimes called a *name server*, is the process that runs on EnGarde awaiting incoming name service requests.

The name server section allows you to specify the name server that will be hosted here. A name server is required for the domain to function properly.

Add Name Server Record
Below you can define namservers for your domain.

Enter your domain in the **Domain Name** field and the name of the name server in the **Name Server** field.

Name Server   `lockbox.guardiandigital.com`

Create

To add the name server simply type it into the *Name Server* field and click on the *Create* button to submit the changes.

Edit/Delete a Name Server

Once you create a new name server you will see it listed below.

Name   **Name Server**
guardiandigital.com.   lockbox.guardiandigital.com.

You can click on the name to edit the record.

Edit Name Server Record
Below you can define namservers for your domain.

Enter your domain in the **Domain Name** field and the name of the name server in the **Name Server** field.

Name Server   `lockbox.guardiandigital.com.`

Save   Delete

To make changes to the record simply edit the field and click the *Save* button. To delete the record click the *Delete* button.

Mail Server

Here you have the ability to set up a mail server for the domain. You can set up more than one server and set the priority level of the server. More detail on doing this will be provided below.

You can define your mail server(s) in the *Mail Server* field. Only one server can be defined at a time. However, you can have more than one mail server per domain with different levels of priority. This provides failover. If a particular mail server is unavailable, DNS will automatically instruct it to use a different mail server.

The order in which the next server is chosen is known as the priority. The lower number the priority, the higher the precedence. In other words, a mail server configured with a priority of 10 will receive mail before one with a priority of 20.

You must complete the *Mail Server* and *Priority* fields. Once you are done, click the *Create* button and the server you just entered in will be displayed at the bottom.

Edit/Delete a Mail Server

Once you have created a mail server it will be listed as shown below.

You can click on the name of the server to bring up the edit screen.

Edit Mail Server Record
Below you can define what machine you want to recieve e-mail for your domain.

Enter your domain in the **Domain Name** field and the machine name in the **Mail Server** field.

| Mail Server | smtp.guardiandigital.com. | Priority | 1 |

Save   Delete

To edit the server simply make necessary changes and click *Save*. Your changes will immediately take effect. To delete the server you can click the *Delete* button.

Edit Zone Parameters

The zone parameters are general settings needed by the zone. You will be presented with a menu of the options with the defaults being displayed. A description of each item is listed below.

Zone Parameters

| Master server | lockbox.guardiandigital.com. |
| Email address | admin@guardiandigital.com |

Save

Master Server   The *Master Server* field contains the address of your master DNS server, also known as a *primary DNS server*. The master server controls the DNS for your zone.

For example, if you own `guardiandigital.com` your master server will be responsible for the hostnames and IP addresses for `guardiandigital.com`.

E-mail Address   The administrative e-mail address responsible for this zone. Generally this is the e-mail address of the system administrator or whomever is responsible for DNS for this zone.

When editing is finished, click the *Save* button to apply the changes.

Edit Zone Options

The zone options are preset to the settings you specified globally in the *Global Options section 4.4.5* on page 99. If you wish to override any global settings you can do so here.

4.4.6   DHCP Server Configuration

DHCP is the Dynamic Host Control Protocol. It allow hosts to obtain a dynamic IP address from a centralized machine. The DHCP server assigns network information for the clients on its network and allows you to control what IP ranges are available for your users.

NOTE: DHCP server is only available if you purchased the Professional Workgroup Suite.

Define Address Ranges

This screen shows all of the address ranges you already have allocated for DHCP. If you would like to define a new range, click on the *Define New Range* link.

After clicking the link you will be presented with the following screen.

All fields must be filled out before you will be able to add this new range. A description of each option is listed below:

Subnet The *DHCP Subnet* is the "network" that the block of IP's is on. For example, if you want to allocate 192.168.1.10 (*Start Address*) through 192.168.1.20 (*End Address*), you would enter 192.168.1.0 here.

Netmask This is the netmask value for the block of IP's you are allocating. A sample netmask is 255.255.255.0. This netmask is sent to the client when they request an address.

Gateway The *DHCP Gateway* is the machine that the client machines need to access to "get to the outside world." This is also referred to as a "default route." When the client machine requests an IP address, this is sent back to them along with the assigned address.

Domain Name This is the domain that the client machines are in. An example value is "inside.xyzcorp.com." This is generally the "domain" portion of the DNS name for the IP address.

DNS Servers These are the DNS servers that the clients should be assigned. A
DNS server is used to resolve names into IP addresses. When the client
requests an IP address, the server will send these DNS servers back along
with the assigned address. You can enter as many DNS servers as you want
here, provided that they are separated with spaces.

Start Address This is the first IP in the range you wish to allocate. If you want to
allocate the range `192.168.1.10` through `192.168.1.20`, you would
enter `192.168.1.10` here.

End Address This is the last IP in the range you wish to allocate. If you want to
allocate the range `192.168.1.10` through `192.168.1.20`, you would
enter `192.168.1.20` here.

When you are done filling out all the entry boxes click the *Create Range* button.

After the new range is created you will be brought back to the previous screen.
You will now see your newly defined range listed here. You have the ability to edit
this range by selecting the *Edit* link associated with the range you wish to edit.

The edit screen is almost identical to the range creation screen with the addition
of a delete button to delete the entire range.

View Current Leases

Whenever a client requests an address via DHCP, the server assigns them the
address and defined a "lease." When the lease expires, the IP is then placed back
into the "pool" of available addresses.

4.4.7   Windows File Sharing

Windows File Sharing allows you to configure your server to host files to Windows based clients. This works by allowing a Windows client to mount a pre-defined directory or share on their own system. Through the WebTool you can define these shares, who has access to them, and what type of access is assigned.

NOTE: This module will only appear if you purchased the Professional Workgroup Suite and chose to install the *Windows File Sharing* package.

This section is broken down into *Global Configuration*, *Machine Management*, *WINS Configuration*, and *Share Configuration* which are discussed below.

Global Configuration

The Global Configuration section allows you to control system-wide settings for *Windows File Sharing*. Here you can configure such options as the workgroup name, machine descriptions, passwords and other items, which are discussed in detail below.

When setting up *Windows File Sharing* computers that will be sharing files with each other will be assigned to a *workgroup* or a *domain*.

A workgroup is used as a way for coworkers to quickly find each other's computers on a network and share files and printers between them.

A domain also contains a collection of computers in a group. They can also browse each other's files and printers, but are required to be authenticated before becoming a member of the domain. This enables the EnGarde Secure Professional server to provide this authentication to the domain members.

Workgroup / Domain If your machine is in a workgroup, then this is the name of the workgroup it should be in.

> If your machine is accepting *Domain Logins*, then this is the name of its domain.

NetBIOS Hostname This is the name the machine will be given when other machines browse the network.

Machine Description This is an informative line that will be displayed when people query for information on this machine.

Local Master? This will set your EnGarde machine to attempt to become the local master browser on your subnet.

Allow Domain Logins? If your EnGarde Secure Professional server is configured as a primary domain controller this will allow other computers to login to the domain of the EnGarde machine.

Share Printers? If this option is set to *yes*, then all of the printers found in the Printer Setup menus in the WebTool will be available to valid users. For more information concerning printers in the WebTool refer to *Section 4.4.10* on page 126.

Interfaces This allows you to enter in the specific IP address you wish to allow Windows File Sharing to accept connections on. Generally you only want this to listen on the internal IP or trusted IP. You can also choose to allow it to listen on all IPs.

Set Administrator Password If your machine is configured as a Domain Master, then you will need an administrative user defined who can authorize new machines to logon to the domain. This is where you set the password for that user.

This option will show up as *Set Administrator Password* if you do not currently have an administrative password set. If you already have one set and you wish to change it, this option will show up as *Change Administrator Password*.

Machine Management

Before a machine can join the domain (if you are accepting domain logins), it must have a machine definition. To define a new machine, go into this section and click on the *Define New Machine* link.

You will then be asked to enter the machine's NetBIOS name into the box. Clicking on *Define Machine* will set complete the machine setup and you can now log into the domain.

Once the machine NetBIOS name has been added it will appear on the main *Machine Management* menu. From here you can edit the entry by clicking on it. You can also delete the entry from within the edit screen or add an additional machine name from that main menu.

WINS Configuration

WINS stands for "Windows Internet Domain Service." It serves the purpose of translating NetBIOS names into IP addresses. If you have the machine set up as a master browser then it will act as a WINS server and will answer any incoming WINS queries.

NOTE: WINS is suitable for environments with no DNS configuration.

To add a new entry click the *New Record* link.

To add static WINS entries to the WINS table you need to enter NetBIOS name and corresponding IP address in this section.

After selecting *Create Record* you will return to the main menu and your entry will appear. To edit this entry simply click on it. You will be returned to a menu similar to the creation menu with the addition of a *Delete Record* button. To delete the entry press this button. Selecting *Save Record* will update any changes you made.

Share Configuration

Share Configuration will allow you to create new shares. When creating a share you define the directory to be shared, the name of the share, who can access it, what groups can access it, or define it as public.

To create a new share click the *Create New Share* button on this menu.

There are three main options that will define who can access your share: *Hosts to Allow*, *Public Share?*, and *Writeable?*. *Hosts to Allow* is a space separated list of IP's or networks that are permitted to connect to this share. This does not define who can access the actual information, it just specifies whether or not a network connection will be established. To allow "all" addresses, simply leave this box blank.

Once a machine is allowed to connect, *Public Share?* specifies weather or not they are allowed to browse the share (read-only). If *Public Share?* is set to *yes*, then all users will be allowed to read the contents of the share. If this is set to *no*, then only *Authorized Users* or *Authorized Groups* will be allowed to browse the share.

Finally, *Writeable?* specifies whether or not to grant worldwide read/write access to the share. If this is set to *yes* then all users who connect will have read/write access. If this is set to *no*, then only *Writeable Users* and *Writeable Groups* will have read/write access to the share.

Share Name This is a label that users will see when browsing.

Directory Enter into this entry box the path to the directory you wish to share.

Share Description This is an informational field the user sees when browsing.

Public Share, Writeable?, *Authorized Users and Groups*, *Writeable Users and Groups* along with *Hosts to Allow* all define access control to a share. The chart below can be used to determine how these options are used to control user access.

Public Shares	Writeable	
No	No	Only Authorized Users/Groups can read the share, on only Writeable Users/Groups can write to the share
Yes	No	Anybody can read the share, but only Writeable Users/Groups can write to it.
No	Yes	Anybody defined in Authorized Users/Groups can both read and write to the share.
Yes	Yes	Anybody can read and write to the share.

When changes are done being made click *Create Share* to create this share.

After the share is created you will be brought back to the main menu. As in *Machine Management* and *WINS Configuration* you have the ability to edit, delete and create new shares at this point.

4.4.8   Network Configuration

Selecting the *Network Configuration* option from the *System Management* section will bring you to the Network Configuration main menu.

The first thing you will see at the top of this menu is the list of interfaces currently installed in your system.

You can edit active interfaces by clicking on the ethernet device link to the left of the interface or edit the virtual address of the device by clicking on its associated *Virtual Address* link to its right. We will discuss more on editing the device later in this section. First we want to create a device. If you click on the *Define New Physical Interface* link you will be brought to a new screen, the *Interface Setup*.

Creating a New Device

Here you can choose to make your new interface use a static IP address that you define, or use DHCP or PPPoE to control the interface.

Static Interface A static interface consists of pre-defined network settings that are restored upon each reboot. If this machine is to be a router, gateway or server, this option is probably for you. Simply select the *Use a static*

address check-box, enter in your IP address and netmask, and save your settings by clicking the *Define Interface* button.

Dynamic Interface (DHCP) DHCP is the Dynamic Host Control Protocol. If you select the *Use DHCP to obtain network settings* check-box then the machine will attempt to contact a remote DHCP server to obtain its network settings. If you are on cable modem or a LAN that uses DHCP to delegate IP addresses, this is probably the option for you.

Dynamic Interface (PPPoE) PPPoE is the Point-to-Point Protocol over Ethernet. If you select the *Use PPPoE to connect to network* check-box then the machine will attempt to connect to the network using the PPPoE protocol. In order to use PPPoE you must have a valid username and password. If you are on a DSL connection then his is probably the option for you.

Selecting '*Yes*' for '*Overwrite DNS Configuration?*' will force this device to use your ISPs DNS servers.

For more information concerning DHCP and PPPoE in regards to a broadband Internet connection refer to *Section 4.4.9* on page 123.

Edit an existing interface

To edit an interface click on the ethernet device you wish to edit. You will see the same menu as if you were creating a new device. Make your changes here, refer to *Creating a New Device* for a definition of each section. When you are done making changes select the *Save Interface* button or you can select *Delete* Interface to remove the selected device from the configuration.

Creating a Virtual Address

To create a virtual interface you can start by clicking on the *'New Virtual Address'* link associated with the device to which you want it bound.

Fill in the IP address you want for this virtual interface and then the netmask. Click the *Define Address* button to apply the changes.

Routing Configuration

In this section you can configure the routing table for the EnGarde Linux system. This is initially configured during the EnGarde installation process but if the physical network was changed since that time or the routing table required updates, this is where it gets done. From here you can define the default route and the static route(s) for the system.

The static route is an explictly defined route. When sending out a packet over the network the static routes will all be searched first. If the packet fails to reach it's destination via the static route(s) it will fall back to the default route, described below.

To add a static route click '*New Static Route*'. A new screen will appear.

Network Enter in the network address of the network this static route is being configured for.

Netmask Enter the address of the netmask for the network defined in the '*Network*' field.

Device Select from the pull-down menu which ethernet device this static route will be configured for.

When all the fields have been correctly filled in clicking '*Define Route*' will create the route and it will now appear on the main '*Routing Configuration*' screen as pictured above.

Below the static route configuration is the '*Default Route*'. The default route is used when a packet fails to reach it's destination via the defined static route(s). If no static routes are defined the defualt route will always be used.

The default route is configured when you install EnGarde on your system. If you wish to make changes modify the appropriate fields.

Gateway You will need to enter in the IP address of the gateway you will be using.

Device This will be the device in your EnGarde system that will be used to access the router. Generally `eth0` is used for this.

When changes are done being made click the '*Save Default Route*' button for changes to take effect.

NOTE: Only configured interfaces will be displayed.

Hostname and DNS Client Configuration

This section will allow you to reconfigure your DNS servers and your hostname, which are configured at installation time. Additionally you can add *Search Domains* from here as well.

Hostname The hostname must be a Fully-Qualified Domain Name. Entering in an incorrect or partial hostname can have serious negative effects on a system. It is also highly recommended not to change the hostname of a production system.

Search Domains Search domains are domains that the system will automatically
search if only a hostname is given. For example, if you specify
"guardiandigital.com" and in your web browser your type "www" in
the address bar the system will know to look for "www.guardiandigital.com"
as well as the other domains you have listed.

Following the *Hostname and Search Domains* configuration is the DNS config-
uration. Here you will see the DNS server(s) that were supplied at install time.
These can be change by typing in new IP addresses into these fields.

NOTE: An IP addres **must** be entered, if a domain name is entered EnGarde
                will not be able to perform DNS lookups.

When changes are done click the *Save Configuration* button to save these changes.

Define Static Host Addresses

When EnGarde is passed a domain name it will use a static host address file to
search first and then DNS to determine the IP address. By entering one or more
Static Host Addresses here you will force the system to use this list first before
searching DNS.

You can only add one at a time. After clicking *Update Hosts* a new entry field will
be available for an additional address.

Restart Networking

Clicking this link will restart the networking on the EnGarde box making effective any changes in the *Network Configuration* section.

Restart Networking   Bring up/down interfaces and make routing changes
                                 active.

NOTE: The default 127.0.0.1 address must not be removed.

4.4.9   Broadband Connectivity

Broadband Internet access has become a common commodity in homes and small businesses as the installation and pricing of cable modems and DSL have been dropping. Below are the requirements for configuring both DHCP and PPPoE devices to work with an EnGarde Secure Linux system.

DHCP Requirements

DHCP in regards to broadband will allow your ethernet device to fetch it's configuration from the DHCP device, such as a cable modem. Configuring an EnGarde system to work with a cable modem can be done easily.

Make certain the cable modem is connected to the ethernet card that is set up for DHCP via the cable supplied with the modem.

Next make certain the ethernet interface connected to the cable modem is configured for DHCP. If this was not done at installation time it can be configured from the WebTool.

After logging into the WebTool select the *System Management* option. Following that select *Network Configuration*. At this point the ethernet interfaces in the EnGarde system will be displayed. A static or PPPoE device can be changed to a DHCP device from here. Refer to *Section 4.4.8* on page 117 for details on how this is done.

The DHCP configuration is now complete. There are now some general configuration requirements that will need to be made. These can be found after the PPPoE Requirements section on page 124.

PPPoE Requirements

PPPoE is short for Point-to-Point Protocol over Ethernet. Point-to-Point Protocol (PPP) is commonly used by analog modems for communication over a phone line. PPPoE allows PPP communications to travel through an Ethernet interface. This method is used primarily for DSL modems.

To configure EnGarde to work with your PPPoE device start by connecting your PPPoE device via the cable supplied with the device, to the ethernet card that will be configured for PPPoE.

If the ethernet device to be used for PPPoE was not configured at installation time the WebTool can be used to accomplish this. Starting in *System Managemen*t select *Network Configuration*. A static or DHCP device can be changed to a PPPoE device from here. Refer to *Section 4.4.8* on page 117 for details on how this is done.

The PPPoE specific configuration is now complete. There are now some general configuration requirements that will need to be made. These can be found in the following section, *Common Configuration Requirements*.

Common Configuration Requirements

Now that DHCP or PPPoE settings have been properly configured, the network needs to be restarted for all the changes to take effect.

To restart the network from the WebTool start in *System Management*, then select *Network Configuration* and click the '*Restart Networking*' link. The network will be restarted when the page refreshes.

At this time you **must** redefine the trusted and untrusted ethernet devices. The system may be more vulnerable to an attack during this time.

To define a trusted and untrusted host go to '*Security*' from the main WebTool menu. Select *Firewall Setup* and then *General Configuration*.

You can select your ethernet devices from the pull-down menu. If you are configuring broadband access you will want to make the DHCP or PPPoE device the untrusted device and the trusted device the device configured for your internal network.

When you have selected the devices click '*Save Configuration*' followed by the '*Restart Firewall*' option above it.

Once the firewall has been restarted it's a good precaution to confirm that DNS is running as expected.

From the main WebTool screen select the '*EnGarde Audit System*' (EAS). At this point select '*Services*' from the pull-down menu, click '*Change Applet*'. A new pull-down menu will appear. Select '*DNS Server*' from this one.

If DNS is running properly the status will be '*Enabled*'. Additionally the DNS service should be '*Enabed*' in the '*At Boot:*' section as well.

More information concerning the usage of the *EnGarde Audit System* can be found in *Section 4.5* on page 134.

4.4.10   Printer Setup

EnGarde Secure Professional allows you to set up your parallel port or USB printer directly through the WebTool. After you have successfully defined the printer connected to your EnGarde server in this section, it will be necessary to install the printer driver supplied by the printer manufacturer on each workstation that wishes to use the printer.

To add a new printer start by clicking the *Define New Printer* link.

After clicking on the link you will be brought to the Printer Setup screen. Here you will need to fill out two options, *Printer Name*, and *Printing Device*.

The *Printer Name* is just a label to give the printer. This name will also be used for the network printer name. Spaces and special characters are not permitted here.

After the *Printer Name* you must select the *Printing Device*. You may choose between a *USB* and *Parallel* printer from the pull-down menu.

Once all the fields have been filled out click the *Create Printer* button and you will see the main screen with your printer listed.

You are now set up to print.

4.4.11   Quota Setup

Quotas are a defined set of rules that limits system resources allocated to each user or a group of users. Resources such as filespace, system processes, memory, etc. can all be limited.

Filesystem Quotas

Filesystem quotas allow you to define how much disk space a particular user or group can use on a given filesystem. When you enter this page you will be shown a listing of each filesystem currently set up. The *User Quotas* and *Group Quotas* will be set to *Enabled* (quotas are being enforced) or *Disabled* (quotas are not being enforced).

In the first section of the filesystem quotas you will see *Define Filesystem Quotas*. Here you will see all of your mounted partitions. By default all quotas are disabled. Since filesystems quotas are disabled you will not see anything listed in the *Existing Filesystem Quotas* section.

To enable filesystem quotas on one of your partitions or to change the partitions options click the *Edit* link associated with the partition you wish to make changes to and you will be brought to the following screen.

Here you have two pull-down menus. Each option is to enable or disable group and/or user quotas. When you have made your selection(s) click the *Save Changes* button to have the new changes take effect.

When finished making changes, click *Save Changes* and you will be returned to the previous screen You will notice now your enabled filesystem quotas are listed for the selected partition in the *Existing Filesystem Quotas* section.

Currently you have enabled quotas for the selected partition but you have not yet defined what these quotas are to be so you will see "No users/groups currently have quotas defined." message. To define a new user quota select the *New User Quota* link, as for groups as well, *New Group Quota* link.

When selecting the New User Quota link you will be brought to the following menu. Here you assign a quota on a per user basis.

User Name Here you can type in the users name or select it from the menu by clicking '...'.

Soft Limit This is a set limit that when reached the user will be informed that they are exceeding there quota but will still allow files to be written.

Hard Limit If the user ignores their soft limit and continues to use disk space they will be denied permission to write anything once they reach this hard limit.

Once changes are finished being made hit the *Create Quota* button for this quota to go into effect.

To create a group quota select the *New Group Quota* link. You will be brought to a screen similar to the *New User Quota* screen but instead of asking for a user name a group name is wanted. This menu works the same way but the quota takes effect for every user of that group.

Resource Limits

The Resource Limits section contains three subsections, *System-wide Limits*, *User Limits* and *Group Limits*. There are all system-wide limits. All the interfaces here work similarly. To edit and existing item select the associated *Edit* link to the right of it and to add a new limit click the *New Limit* button associated with it.

System-wide Limits

Domain	Type	Item	Value	
Any	hard	Maximum Core Size (KB)	0	[Edit]
Any	hard	Maximum Number of Processes	160	[Edit]
Any	hard	Maximum File Size (KB)	40000	[Edit]
Any	hard	Maximum Number of Logins	5	[Edit]

[New System-wide Limit]

User Limits

Username	Type	Item	Value	
mysql	hard	Maximum File Size (KB)	10000000	[Edit]

[New User Limit]

Group Limits

Group	Type	Item	Value	
users	hard	Maximum File Size (KB)	15000	[Edit]

[New Group Limit]

System-wide Limits

All the limits set in here are generic limits that effect everything that is not con-
trolled by the root user. You have three options from this menu, if your limit is
soft or hard, what kind of limit it will be and the value of the limit. Each item is
broken down below.

Resource Limit Maintenance

Domain	Any	Type	Hard ▭
Item	Maximum File Size (KB) ▭	Value	2500

Create

Type This type allows you to choose between a *Soft* and *Hard* limit. A soft
limit informs the user that they have exceeded their quota while a hard limit
"cuts" the user off, preventing them from using any more resources.

Item This is a pull-down list of items that describe how this limit will behave:

- **Maximum core size(Kb)** - This limits the size of a core file. A core
  file is a file that a program will write to the system when that program
  crashes. The developer can then take this core file and use it for de-
  bugging the program. If the system is not used for developement it
  should be set to `200`.

- **Maximum file size (Kb)** - This is the maximum size a single file is allowed to be. This option is desireable for enforcing e-mail mailbox limits.
- **Maximum Logins** - This controls the maximum number of simultaneous logins
- **Maximum Number of Open File**s - This limits the total number of open files on the system. An open file is any file with its flag set to open.
- **Maximum Number of Processes** - This will limit the total number of current running processes on the system.
- **Maximum RSS Size (Kb)** - This specified the total amount of physical memory used, not counting pages swapped out

Value This is the numerical value associated with the item. For example, if you chose Maximum file size (Kb) then a value of `250` would be `250Kb` or if you selected Maximum Logins then a value of `5` would denote a maximum of 5 logins.

User Limits

The User Limits here will allow you to set what was optional in the previous *System-wide Limits* section on a per user basis.

Fill out each entry box first with the name of the user, followed by the Soft/Hard option, item and value as described in the *System-wide Limits* section prior to this.

Group Limits

The *Group Limits* allows you to set everything like you did in the *User Limits* section but the changes effect an entire group instead of a single user. Refer to the *System-wide Limits* section for a description of each field.

4.4.12   Change System Time

This section allows you to change the current system time, or synchronize it with an Internet or designated local time server.

Changing the time is controlled by pull down menus. Select the current time and hit *Set System Time* for the changes to take effect. Normally, system time will be accurately controlled with the network time services and manually setting it is not necessary.

It is also possible to configure EnGarde to use Internet time servers to set its time.

You have three fields to fill in the hostnames of the time servers. EnGarde will use all three servers to synchronize its time. Keeping accurate system time is extremely important. You have to enter hostnames in here. IP addresses are not allowed.

Setup Time Servers

This section allows you to setup which servers you would like to use as time servers. For more information on this, please visit http://www.ntp.org. Your system will sync itself with these systems often, using the three servers to assure accuracy. Your current time servers are listed below the text boxes.

A list of servers can be found here. Please find the three geographically closest to you and enter them in the text boxes below. (Note: These should be hostnames, **not** IP addresses.)

Server One   **Server Two**   **Server Three**

rrapin.csc.ncsu.edu   tock.usno.navy.mil   bonehed.lcs.mit.edu

terrapin.csc.ncsu.edu   **tock.usno.navy.mil**   **bonehed.lcs.mit.edu**

Setup Servers

4.5 EnGarde Auditing System (EAS)

Auditing is the process by which EnGarde lets you know what's going on with both users and processes on the system, as well as how it is currently performing. Information must be checked for internal consistency and for consistency with other criteria. The EnGarde Auditing System provides an audit trail that enables administrators to reconstruct later who did what, in case it is suspected there may be a system anomoly.

The *EnGarde Auditing System* (EAS) allows recent system logs, Web logs, and graphs of network and system events to be viewed. Additionally, the system can be shut down or restarted from here as well.

To select different options click on the pull-down menu, select the option by clicking on it and click the '*Change Applets*' button.

4.5.1 System Graphs

The *System Graphs* section will display several graphs of different system statistics. By clicking on a graph a daily, weekly, monthly and yearly breakdown will be displayed.

Information such as ethernet usage, memory usage, CPU usage and CPU temperature are displayed in these graphs. Below is a sample graph of ethernet usage.

4.5.2   Services

The *Services* section allows you to choose from the different services on your server from a pull-down menu.

After selecting a service from the pull-down menu you will be represented with the current status of the service, whether the service is being started at boot time and the ability to toggle these two options.

Below these two options you will additionally see the most recent logs generated from the selected service.

4.5.3   Website Logs

The *Website Logs* will display the most recent logs from a selected Web site hosted on the EnGarde server. To choose which of your Web sites you wish to view logs from select one from the pull down menu.

4.5.4   System Reports

System Reports are run nightly and contain information on the currently running system. Such things as free memory, open port, current connections, disk usage, e-mail statistics, DNS statistics and others can be found in this report.

To choose a report for a specific day select it from the pull-down menu and click the '*View'* button.

The report for the selected date will then appear in the browsers window.

```
+---------------------------------------------------------------+
| bluehen.inside.guardiandigital.com                            |
| EnGarde Secure Linux, version 1.1 (Balestra)                  |
+---------------------------------------------------------------+

 Report Time:     Wed May 15 14:00:01 EDT 2002

 System Uptime:  1:23
 Load Average:   0.00, 0.01, 0.00
 Kernel Version: 2.2.20-1.2.23ipsec

#
# Network Device Information
#

 Interface   Address/Mask        RX Packets %err  TX Packets %err
 --------------------------------------------------------------
 eth0        192.168.1.7/24          15895   0%       10225   0%
 eth0:1      192.168.1.6/24            n/a              n/a
 ipsec0      192.168.1.7/24            n/a              n/a
 lo          127.0.0.1/8
```

4.5.5   Process Information

Process Information contains a list of the current running processes on the system. You can choose to arrange them by '*User*', '*%CPU*', or '*%Memory*' by clicking on the link at the top of the process list.

```
Sort By: [ User ] [ %CPU ] [ %Memory ]

User   % CPU  % Memory  Command
root   0.0    0.1       init [3]
root   0.0    0.0       [kflushd]
root   0.0    0.0       [kupdate]
root   0.0    0.0       [kswapd]
root   0.0    0.0       [keventd]
root   0.0    0.0       [mdrecoveryd]
root   0.0    0.2       /sbin/syslog-ng --cfgfile=/etc/syslog-ng.conf
root   0.0    0.3       klogd -c 1
root   0.0    0.2       crond
root   0.0    0.3       sh /usr/lib/ipsec/_plutorun --debug none --uniqueids ye...
root   0.0    0.1       logger -p daemon.error -t ipsec__plutorun
root   0.0    0.3       sh /usr/lib/ipsec/_plutorun --debug none --uniqueids ye...
root   0.0    0.3       sh /usr/lib/ipsec/_plutoload --load %search --start %se...
root   0.0    0.3       /usr/lib/ipsec/pluto --nofork --debug-none --uniqueids
root   0.0    0.3       xinetd -reuse -stayalive
```

4.5.6   System Control

System Control gives you two options, *Reboot System* and *Shutdown System*. You will need to check the check-box associated to the option you wish to use before clicking the button. This is done to prevent accidentally clicking a button and bringing down the system.

4.5.7   Edit Configuration

The *EnGarde Auditing System* gives you full control over how the system information is visually displayed. Here you can change such options as the number of lines in a log to display, refresh time and window size. See below for a detailed list of each option.

When you are finished making your changes click '*Save Changes*' for the new changes to take effect.

EAS Window Width This will set the width of the pop-up window the EAS uses.

EAS Window Height This will set the height of the pop-up window the EAS uses.

EAS Window Scrollbars? Selecting '*No*' will remove all the scrollbars from the pop-up browser windows.

EAS_Window_Menubar? Select 'No' will remove the menubar from the pop-up browser windows.

EAS Refresh Time Each pop-up window will be refreshed after X seconds. Set how often you wish to have your windows refresh.

EAS Display Lines This is the number of lines the EAS applets will display. Its meaning varies from applet to applet. For example, in the Services applet it defines how many log lines will be displayed and in the Process applet it defines how many processes will be displayed.

EAS Truncate Length This is the number of characters (on a line) that will be displayed before the line is truncated in the interest of display. This is used to control wrapped caused by long lines in the pop-up.

4.6   Security

EnGarde Secure Professional includes all necessary security settings pre-configured.
They are optimally set for the highest level of security without hindering the usage
of EnGarde. This section will let you configure some of these security settings to
adapt to possible system changes you may make over time. From here you have
the ability to manage certificates, configure SSL encryption, IP access control,
customize your console login banner, configure host intrusion detection, gateway
firewalling and virtual private networking.

4.6.1   Change WebTool Password

You can change your administrative WebTool password here. You need to enter
it in twice to avoid typing errors. We recommend a password no shorter than six
characters. Mixing letters and numbers is a good idea and avoid full words. See
LinuxSecurity.com for tips on choosing a secure password.

4.6.2   Change Administrator E-Mail Address

The administrators address can be entered here to receive a daily summary of important log information and security alerts.

The Daily Summary

The daily summary is e-mailed out every night at ten minutes past twelve. The contents will look something like this sample daily summary e-mail:

```
Log Summary for 10/3/2000

*** Log summary for system logins ***
Total number of:
 - root logins via su                  - 0
 - SSH sessions opened                 - 5
 - console logins                      - 0
```

```
*** Log summary for GD WebTool logins ***
Total number of:
 - successful administrator logins      - 16
 - failed logins                        - 4


This has been e-mailed to : nick@guardiandigital.com

End of summary for 10/3/2000
```

Depending on your system configuration and installed packages, you may receive more or less information in this summary.

Security Alerts

For servers that have the LIDS host intrusion detection service enabled, and someone tries to disable it, but gives an incorrect password three times in a row in under a one minute interval, an e-mail will be sent to the administrator whose address was specified in the *Change Administrator E-Mail Address* section.

NOTE: Chances are you can safely ignore this section. If you are uncertain of what to do should this event arise, contact Guardian Digital for further assistance and we will be glad to help.

The e-mail will contain instructions on how to handle the situation. It will look similar to the example below:

```
A password to disable the host intrusion monitor was en-
tered three (3) times incorrectly. This could be an er-
ror of the system administrator or it could be some-
one attempting to gain unauthorized access.

We suggest checking in to this matter as soon as possi-
ble. To check if the host intrusion monitor is prop-
erly running login to your Lockbox as the root user. In-
structions on this can be found in Section 6 of the docu-
mentation, and type:
```

```
lidsadm -r

This will return the current running status of the intru-
sion monitor. If the monitor is not run-
ning you should turn it back on. Do this by typing:

lidsadm -S -- +LIDS_GLOBAL

It will prompt you for your host intrusion monitor pass-
word. Once the password is correctly entered the intru-
sion monitor will be en-
abled. You can scan the logs through the GD WebTool for more de-
tailed information. You can also read more on the intru-
sion monitor in Section 9 of your of your manual.
```

This error will only occur under the following conditions:

- A wrong password is entered in three times in a row to disable LIDS

- A wrong password is entered in three times in a row to enable LIDS

- A wrong password is entered in three times in a row to reload the LIDS configuration

What this means is that either a user with root access accidently entered in the password wrong three times in a row or an unauthorized user has attempted to gain access.

If you only use the GD WebTool to administer your EnGarde system you should rarely see this message.

In the event of this e-mail, you are welcome to contact Guardian Digital for further assistance. Read *Section* 1.4 on page 8 on how to contact Guardian Digital.

4.6.3   Edit Login Banner

This allows you to alter the login banner the user sees when they connect to the system or login from the console. Just type in plaintext and hit *save* when finished. We recommend putting in a warning/disclaimer about illegally accessing the system. It may be necessary to consult your security or legal department.

4.6.4   WebTool Access Control

This section allows you to control what IP addresses have access to the GD WebTool. You should allow as minimum as possible. You can enter the IP addresses in a list, entering a new line after each entry.

Choosing the *Allow from all addresses* option can place your system at the greatest security risk.

4.6.5   System Access Control

This works similar to the *WebTool Access Control* section except these rules apply system-wide.

Entering an IP address in the given *IP Address* field will allow that IP Address to connect to EnGarde using the selected service. Checking the '*Allow all addresses*' check-box will allow any and all IP Address to access the selected service. Examples are given above the IP Address field.

Once you have that typed in click the *Add Host* button and your new settings will appear below once the screen refreshes.

4.6.6   Secure E-Mail Client Setup

EnGarde Secure Professional supports both Secure IMAP (simap) and Secure POP3 (spop3). Here you can configure which interfaces each service can listen on and configure your certificates for each service.

Secure IMAP and POP3

Both the Secure IMAP and Secure POP3 interfaces allow you to configure which network interface(s) you want each service to listen on. By leaving the entry box blank, the service will listen on all network interfaces. To select a specific network interface you can type in the IP address of the network interface or click the ' . . .' button for a list of available interfaces.

Edit Certificate

Both services come with their own default certificate issued by Guardian Digital. You change this certificate as you wish through the *Edit Certificate* interface.

Authority Name This is the name of the host the certificate will be used on. This name must match your FQDN for SIMAP and SPOP3 to both work properly.

E-Mail Address This is the authoritive contact. This can be an individuals address in charge of the address or the system administrator.

Organization The organization is the name of the company or individual who will own the certificate.

Department This is a sub-category of the company name. You should enter in the name of the department within the organization that has control over this certificate.

City This is the city that the physical server resides in.

State or Province This is the state or province in which the city, from the above definition, resides in.

Country The country entry box requires a two letter code to designate your country.

4.6.7   Tripwire Maintenance

Tripwire is an open source security tool copyrighted by Tripwire Security, Inc. and customized for EnGarde by Guardian Digital. It that monitors changes in

file attributes and will raise an alert via an e-mail to the system administrator concerning file changes that should not have taken place.

When you first visit to the Tripwire Maintenance section there will be instructions for initializing the Tripwire configuration. This must be done before you can access the WebTool's Tripwire module.

Tripwire Initial Configuration

Before you use this module, you must first set up Tripwire on your machine. Unfortunately, this cannot be done through the WebTool at this time so it must be done from a shell. Below are the steps you will need to perform. Be sure that these commands are executed as the 'root' user.

```
[root@machine/root]# /etc/tripwire/twinstall.sh
[root@machine/root]# tripwire --init
```

After you execute the 'twinstall.sh' script, you will be prompted for a *site keyfile passphrase* and a *local keyfile passphrase*. These passphrases should follow the guidelines outlined in the *Initial Configuration* section of this manual on page45.

After the keys are generated, you will be prompted for your *site passphrase* two times, as Tripwire signs its configuration files with this key to ensure data integrity.

When that script is done, you can run the second command to initialize your database. You will be prompted for your *local passphrase* when initializing the database. If you see *No such file or directory* warnings do not be alarmed. The configuration file provided in EnGarde covers a stock installation, with all services running. If you have some services disabled then Tripwire will generate these harmless warnings. These warning can be addressed in the *Tripwire Maintenance* section of the WebTool.

The first time Tripwire is run, a reference database will be created that reflects the normal operating state. Variations from this reference database will require intervention to include these changes to reflect this new state. It is therefore recommended Tripwire be initialized only after your system is fully configured and before being connected to a network to minimize the potential for variation.

It is recommended Tripwire be started *after* your system is fully setup. The administrator will be notified of any changes from the point Tripwire is started and could become a hassle if the system is still being configured.

Once these steps are performed, you can come back to this section to use the WebTool module.

Tripwire Maintenance

Once the steps to initially configure Tripwire listed above have been completed you will see the following screen upon returning to the *Tripwire Maintenance* section.

Define Administrator

Tripwire sends a daily report informing you of any system changes. To change who gets the message, type in the e-mail address of the person to receive the reports and enter in the passphrase you used to set up Tripwire.

NOTE: This process will take about 4 minutes. Please do not click the *Stop* button or interrupt the process.

Schedule Tripwire

Tripwire is scheduled by default to run at midnight everyday. Using the pull down menus you can change how often and when Tripwire is run.

Generate & View Reports

You can force Tripwire to create a report by selecting *Generate Report*. After it has finished generating a report you can get a list of all the recently generated reports be selecting *View Reports*.

Selecting a listed report will display the report to you with the option to delete the selected report at the bottom of the report.

Update Database

When you select the *Update Database* option Tripwire will create a list of all the files that have changed and will display them to you along with a check-box next to each one.

To add an item to the database unselect the check-box. Once all changes have been made enter in the passphrase and select *Update Database*.

4.6.8   Firewall Setup

EnGarde Secure Professional allows you to configure global firewall settings and set up port forwarding rules. The firewall security policy configured with EnGarde by Guardian Digital, combined with the additional security measures included with EnGarde, provide a robust firewall configuration for most environments. A description of each menu and the items contained within it are explained below.

General Configuration

In this section you are asked to define the *Trusted Interface* and the *Untrusted Interface*. Generally, the "Trusted" interface is the one that is connected to your internal network and the "Untrusted" interface is the one that is connected directly to the Internet.

The firewall rules that are in effect block all incoming Windows Networking, DHCP and syslog communication from the outside as well as translate external requests for services by internal workstations using Network Address Translation.

Firewall Modules

The *Firewall Modules* are a collection of IP masquerading modules to allow protocols such as FTP, IRC, PPTP and a few others to be transferred through the firewall. You can enable these modules by selecting *Enabled* from the pull-down menu.

NOTE: If you have only one network card you should set these to *Disabled*.

Firewall Status

The *Firewall Status* section will show you the current running status of the fire-wall, either enabled or disabled and allow you to shutdown, turn on or restart the firewall.

You can toggle the firewall on and off by clicking the *click here* link next to *Disable Firewall*.

Clicking the *click here* link to the right of *Restart Firewall* will restart the firewall.

Configuration Options

The *Configuration Options* allows you to tell the EnGarde Secure Professional server which network interface is your *Trusted Interface* and which one is your *Untrusted Interface*.

Generally the external interface is the *Untrusted Interface* and the internal interface is the *Trusted Interface*.

NOTE: This section will not appear if only one network interface is present in the system.

Configure Port Forwarding

Port forwarding is a method for forwarding requests for service to a server that would otherwise not be reachable from the external network. This enables an organization with a single publically-accessible IP address to potentially forward services such as HTTP and SMTP to servers located within their internal network.

The diagram in *Figure 4* on page 155 describes a typical scenario where an En-Garde Secure Professional server is configured to forward SMTP requests to an organization's internal mail server using the publically-accessible IP address assigned to the EnGarde Secure Professional server itself.

The following steps correspond to the sections in the diagram as data traverses from the workstation on the Internet to the internal server and back to the workstation.

(A) The end-user on the Internet makes a request for a webpage.

(B) The request passes through the Internet and makes its way to your EnGarde server. The EnGarde server will evaluate what type of request it is and determine how to handle it based on the rules defined by the EnGarde administrator in this section.

(C) In this example it will forward the request to the web server located on the internal network.

(D) The web server will handle the request and send the results back to the EnGarde server.

(E) The EnGarde server at this point will forward the results back out to the Internet and to the end-user's PC.

EnGarde gives you the ability to set up port forwarding directly through the WebTool. Here you can define what service requests addressed to the external interface of the firewall will be passed on to servers on the internal network.

When you first visit this section you will not see any rules listed.

To add a rule select the *Define New Rule* link. You will be presented with the following screen.

Here you get to configure and create the new rule. You have the following fields to fill out:

Protocol Select the protocol, TCP or UDP you wish to use for this rule. This should correspond to the protocol used by the port selected.

Figure 4: Port Forwarding Example

Port The ports are listed by their associated services, with the port in parenthesis. Select which service you wish to forward.

Local Address Select the local address (the address on this machine) that you wish to forward from. This will generally be an external interface of the firewall.

Remote Access This is the address you will be forwarding to. This will generally be a server on internal network of the firewall.

The example above describes how to forward SMTP (port 25) on IP address `209.10.240.72` to the SMTP port on IP address `192.168.100.100` on the internal side of the EnGarde Secure Professional server. All requests for SMTP from the outside world to `209.10.240.72` will be forwarded to the internal server on IP address `192.168.100.100`.

NOTE: It is important to note that when port forwarding from the external interface of your EnGarde Secure Professional server to a server located on the internal network, DNS services may need to be configured differently.

Most organizations configure one domain that is accessed by the public and corresponds to the public IP address assigned to the external interface of the EnGarde Secure Professional server.

Internal users accessing the internal server then use a different domain since the server is local to them and corresponds to a local IP address not reachable by Internet users.

This avoids the problem that arises as a result of users attempting to reach the service that is forwarded by the EnGarde Secure Professional server back to the server that is already local to them.

Once everything has been filled out select *Define Rule*. You will be brought back to the main screen and it will display the new rule that was just created.

Protocol	Local Host / Port	Remost Host / Port	
tcp	209.10.240.72 (ssh)	192.168.100.100 (ssh)	[Edit]
			[Define New Rule]

At this point you can create more rules or edit existing rules by selecting *Edit* next to the associated rule.

The *Edit Rule* menu is the same as the *Create Rule* menu except with a button to delete the rule. Delete the rule by simply clicking the *Delete Rule* button.

4.6.9   Virtual Private Networking

EnGarde Secure Professional uses the PPTP protocol to create virtual private networks. This protocol is used by Microsoft clients to create a VPN, or a secure private communications channel between two computers. In the *PPTP Setup* you can configure PPTP options and define new users.

NOTE: This module will only appear if you purchased the Professional Workgroup Suite and chose to install the PPTP package.

General Configuration

In this section are the general configuration options that apply to all connections such as the local IP address to use, the address ranges to issue to remote clients, and what address the daemon should listen for connections on can be configured.

Verbose Debugging Messages If this option is enabled PPTP will produce very
verbose log messages in /var/log/messages. This should be dis-
abled under normal circumstances. If you are having trouble with PPTP
you should enable this option and see what messages are showing up in
/var/log/messages.

Local IP Address This is the IP address that the local PPPTP daemon will bind
to. This should be the IP, or virtual IP address of the machine that your
PPPTP connection will be coming from.

Remote IP Address These are the ranges of IP addresses that the PPTP daemon
will hand out to connecting clients.

You can specify single IP addresses separated by commas or you can specify
ranges, or both. For example:

192.168.0.234,192.168.0.245-249,192.168.0.254

IMPORTANT RESTRICTIONS:

1. No spaces are permitted between commas or within addresses.

2. No shortcuts in ranges! ie. 234-8 does not mean 234 to 238, you must type 234-238 if you mean this.

3. You MUST give at least one remote IP for each simultaneous client.

Address to Listen On This is the address off an interface on the machine that will listen for connections. Leave this blank to allow all interfaces to listen.

Local WINS Server This is the IP address of your WINS server. If you setup your EnGarde machine as a Windows File Sharing server then the IP address of the EnGarde machine can be used.

40-bit Encryption This specifies whether the PPTP daemon should use 40-bit RC4 encryption / compression for the key. 40bit encryption will be used if the client does not support 128bit encryption, or if 128bit encryption is disabled. It is recommended this option remains enabled.

128-bit Encryption This specifies whether the PPTP daemon should use 128-bit RC4 encryption / compression for the key. This will use 128bit encryption as opposed to 40bit encryption if the client supports it.

Stateless Encryption This specifies whether the PPTP daemon should use stateless encryption. It is highly recommended you have this feature enabled. Stateless encryption will randomly change the key during the session which in turn greatly increases security. Without this enabled the same key is used for the entire session.

Edit User

Here you can define, edit and delete PPTP users. This interface will list all the users once they have been created. To create a new user click on the *Create New User* link.

At the Create New User screen you assign the user a user name and password. When you are done click *Create User* and you will be returned to the main menu.

Username This is the username required to establish the VPN.

> It may be necessary to specify the users workgroup in some cases (and certain Windows configurations). The syntax for this is:
>
> ```
> workgroup\\username
> ```

Password This is the users password. Please note that this is kept in cleartext on the machine.

Once you are returned to the main menu the user will appear there. You can now add another user or edit a user by clicking on their username. From the edit menu you can delete the user.

4.7   System Backup

Backing up your system is one of the most crucial roles of system administration. The system backup section allows you to completely backup all characteristics of your system. You can backup configuration files, user home directories, define your own backups, or backup the whole system from here.

4.7.1   System Backup Configuration

The *System Backup Configuration* menu contains general configuration options and your backup options.

General Configuration

In the *General Configuration* section you have to choose your method of backing up. EnGarde supports SCSI and IDE tape drives for backup and will also allow you to backup to a file located on your hard drive.

There are also two other options in the menu, *Overwrite Newer Files?* and *Rewind only?*.

Overwite Newer Files is only applicable if you set if your *Backup Method* is *Backup to File*. If *Overwrite Newer Files?* is set to *Yes* files being restored will overwrite files on the system newer than ones that already exist.

Rewind only? affects only tape backups. If *Rewind only?* is set to *Yes* it will rewind the tape to the beginning when making a backup instead of erasing a tape. This is done since erasing a tape could possibly take hours.

Select which method you wish to use from the pull-down menu and use the *Save Configuration* button to confirm the changes.

Define Named Backup

The WebTool comes with a list of predefined backups. These are all disabled by default. To create new ones select *Define New Named Backup*. Click on this button and you will be brought to a new menu.

The *Exclude Patterns* field is the only optional field in this menu.

Backup Description You will need to give your backup a descriptive name. In
the example above we will be backing up our database files for MySQL, so
it was named *MySQL Backup*.

Directory This is the directory path containing the contents of what you want
backed up. In the example we are backing up all the database material so
we pointed it to the top level database directory. It will backup recursively.

Exclude Patterns Using standard wild card flags and regular expressions you can
choose files not to be backed up. In the example we didn't want to backup
the error files, so we entered in `*.err`. All files ending in `.err` will be
excluded from this backup.

Backup Schedule This pull-down menu contains four options concerning when
you want this backup executed, *Never*, *Daily*, *Weekly* and *Monthly*. Select-
ing *Never* will disable this backup, but it will not delete it.

Backup Level The *Backup Level* will give you two options, full and incremental.
Full will backup every file while incremental will backup only files that
have changed since the last backup.

Once you have everything filled in hit the *Define Named Backup* button and you
will be brought back to the main screen with your new named backup now listed.

To edit one of the predefined backups or to edit a newly created one you can select the *Edit* link associated with the backup. This will bring you to a screen almost identical to the *Define Named Backup* screen and will give you the option to update or delete the named backup. You can also enable a predefined named backup here.

4.7.2   Perform Tape/Directory Maintenance

The WebTool offers the ability to help maintain your backups. If you configured your backups to use tape then you will see the option to initialize the tape, which consists of erasing it, resetting it and setting the system up for use with a blank tape.

Otherwise, if you selected to backup to the hard drive you will have the option to initialize the backups on there by clearing out old backups and initializing the directories so they are ready to accept new backups.

4.7.3   Create a New Backup

Creating a new backup will allow you to run one of your predefined named back-ups immediately. Don't confuse this with the ability to create a new type of backup. When you select the *Create a New Backup* link you will be brought to a new menu.

You will have the option to choose a backup to perform. When you made your selection hit the *Select* button. You will then be prompted to choose between an incremental and full backup.

Once you have done this the backup will proceed and after everything is finished a summary screen will be displayed showing the size of the final archive and what files are contained within it and your backup is complete.

4.7.4 Restore a Backup

If you find the need to restore one of your old backups you can quite easily accomplish this through this interface. When selecting to *Restore a Backup* you will be brought to a menu listing all the named backups, almost an identical menu as the *Create a New Backup* menu.

Select which named backup you want and choose *Select.*

At this point you will be brought to another menu listing all of the backups listed under this named backup. To the right of each named backup is the *Toggle List* option. Clicking on this link will display a list of all directories contained in the backup set. You can choose to restore only portions of the backup or the entire backup. Select which one specifically you want to use to restore with select *Restore Backup*.

All the data in the backup will overwrite all current data so you are asked to confirm your decision after selecting *Restore Backup*.

After confirming your decision you will see a screen giving you a summary of what files were restored, similar to the summary screen in *Create a New Backup*.

4.7.5   View Changes Since Backup

The *View Changes Since Backup* option will allow you to compare the current files on the system against a backup of your choice.

When you first click on *View Changes Since Backup* you will see a screen similar to *Create a New Backup*. Select which named backup your backup is located under, then hit *Select*. You will be viewing a list of all the backups you have made in this named backup.

After choosing which one you want to use to compare with hit the *Diff Backup* button. Because this could possibly put a heavy load on the system you are asked to confirm your decision.

Once you hit *Really Compare* the process will begin. Upon completion you will see a summary screen, similar to when you create a successful backup, listing all the changed files.

4.8   Secure Manager

As discussed earlier the administrator has the ability to change a users password from the WebTool. To increase security, the WebTool does not allow any user but the administrator access to those sections of the WebTool. To allow a user to change their own password themselves, a separate URL is provided. By going to:

```
https://myserver.com:1022
```

NOTE: The address is very similar to the regular WebTool but notice the port you are connecting to. The port 1023 is used for the WebTool, while 1022 is the user password utility, as in the example above.

If the default Guardian Digital certificate still remains on the system the user will be prompted to accept it. Instructions on accepting a certificate can be found in *Appendix E* on page 289.

Once the user successfully logs in to the system using their own login name and password, they will be given the options to either change their password or their secure shell (SSH) key.

4.8.1   Change System Password

In this section a user can change their system password. The old password must first be entered followed up by the new password twice. If both new passwords match the user will be logged out and the password will be updated.

Clicking the *Abort and Log Out* button will cancel this operation.

4.8.2   Secure Shell Key Management

Here the user has the ability to create or upload their own public key to the En-Garde Secure Professional server so that they may be able to SSH into the server. For more information on what SSH is and how to use it in a Windows and Unix environment refer to *Section 6* on page 179.

The main menu here is broken down into three sections, *Keys in your Keyring*, *Upload a Public Key* and *Generate a New Keypair*.

Keys in Your Keyring

This section is only for viewing current keys and deleting them. When you first visit this section there will be nothing listed here since there are no keys in the system.

If you have already uploaded or generated a key it will be visible from here. Clicking on the *[Remove]* link will remove it from the server.

Upload a Public Key

Here a user can upload a public key that you have previously generated. You can type in the path to the key or use the *Browse* button to find it.

Once the path to the key is in the entry box, *Upload Key* can be clicked to upload the key to the server. Once it is uploaded you will see it listed in the *Keys in Your Keyring* section.

Generate a New Keypair

Here a user can create a new keypair. This will create the keypair on the EnGarde Secure Professional system and give the user a copy of the key so they may login remotely.

Filename This filename is the name that will be used to store your private and public key on the EnGarde Secure Professional server. They filename must be alphanumeric.

Description This description is displayed when trying to connect to the EnGarde Secure Professional server using this key.

Passphrase The passphrase is used to authenticate the user and works similar to a password. It will need to be entered twice to check for typing mistakes.

Once all the fields have been filled out click the *Generate Pairkey* button to create the keys. You will then be prompted to download you key to your PC.

Clicking the *Click Here to Download Private Key* button will prompt your browser to download the key. A default filename is given that corresponds to the server and user name, this can be changed.

At this point the new key will be listed in the *Keys in Your Keyring* section.

5 GUARDIAN DIGITAL SECURE NETWORK

Whether you're a small organization new to the Internet world, or a large organization with dozens of EnGarde Secure Professional servers, your security needs are just as important. A security system that is out of date leaves you more susceptible to cybervandals. Maintaining system security, keeping up to date with the latest software improvements, and obtaining access to technical support has been difficult, until now.

Guardian Digital's Secure Network is a means to keep your systems updated while at the same time receiving authoritative advice, information, and additional services from the experts. As you focus on building your Internet presence, Guardian Digital experts focus on assuring you are protected from cybervandals and developing system improvements. Guardian Digital has a dedicated group of security experts that both monitor security sources on a constant basis to identify potential vulnerabilities as well as actively audit the core components of EnGarde, improving the overall security it provides.

Guardian Digital Secure Network is the least expensive way to add dedicated security experts to your staff focused on keeping your systems secure and up to date. This vigilant approach to system security and management is the most effective means to protect your corporate assets and remain up-to-date.

Protect your investment and lower support costs, while at the same time improving the security and functionality of your EnGarde servers. The Guardian Digital Secure Network is available as a monthly or annual subscription.

5.1   Running Guardian Digital Secure Network

To start the Guardian Digital Secure Network select the *Guardian Digital Secure Network* icon from the main menu. You will be brought to the main Guardian Digital Secure Network menu. From here general configuration changes can be made, packages installed from CD media and updated packages downloaded from Guardian Digital.

The purchase of EnGarde Secure Professional includes a trial subscription to the Guardian Digital Secure Network. To take advantage of the features included in the Guardian Digital Secure Network, you first must activate your subscription by visiting:

```
https://www.GuardianDigital.com/register
```

You will be issued an activation password which must be entered into the Guardian Digital Secure Network configuration, detailed below.

To purchase a subscription to the Guardian Digital Secure Network beyond the trial period, visit the Guardian Digital online store by clicking on the *Guardian Digital Store* icon from the WebTool.

The Guardian Digital Secure Network is authorized for use on one EnGarde Secure Professional installation. A Guardian Digital Secure Network subscription must be purchased for each copy of EnGarde Secure Professional installed on your network.

5.1.1   General Configuration

This section allows you to control a few global functions of the Guardian Digital Secure Network. Here you can select to use an advanced mode and enter in

the account number and password, supplied by Guardian Digital for use with the
Update Agent.

Auto-Check Agent Selections If this is set to *Enabled*, then updated package (in
the Update Agent) will be auto-selected for retrieval. If this is not set then
you will have to check each package individually.

Advanced View If this is set to *Enabled*, then dependancy information will be
show in the update agent.

Activation Code This is the number assigned to you from Guardian Digital when
you registered your copy of EnGarde Secure Professional. This allows you
access to the *Update Agent* so that you can update your EnGarde Secure
Professional with the latest packages directly from Guardian Digital.

Account Password This is the assigned password you also received when regis-
tering to be used along with your *Account Number*.

5.1.2   Install from Local Media

The *Install from Local Media* section will allow packages to be installed from
CD-ROM media supplied by Guardian Digital. If you purchased the Professional
Workgroup Suite you would install from here.

To install packages from a CD insert the CD into the CD-ROM drive located in
the EnGarde Secure Professional server. From the main *Guardian Digital Secure
Network* menu select the *Install from Local Media* link. This will prompt EnGarde

to mount the CD-ROM and evaluate its contents. This may take a few moments as EnGarde gathers information about the packages.

Once all the information is gathered you will be presented with a list of packages, descriptions and an option to install them. This will only display packages that are *not* installed on the system.

Select which packages to install by clicking the *Yes* button located next to it. When all selections have been made click the *Install Packages* button. After clicking the *Install Packages* button the packages will begin to install. This will take a few moments and your browser will wait for it to complete. Do **not** hit *stop*, *back* or *reload* in your browser during this process or the packages will not install correctly.

When the packages have finished being installed a screen displaying the packages that were installed will appear. Next to each package will be a link to another portion of the WebTool that is used to configure that package, if available. Using this link will open a new browser window.

5.1.3   Run the Update Agent

The *Update Agent* will contact Guardian Digital servers and over a secure connection determine which packages can be updated. When a list has been determined the screen will display a list of all packages that are newer than what is currently on your EnGarde Secure Professional system.

The screen will show *Severity* of the update, a link to the *Advisory* web page, the *Installed Version* currently on the server, the new *Available Version* and if all the *Dependancies* are met will all be listed.

To select a package to download click the check-box labeled *Retrieve*. When finished making the selections click the *Retrieve Packages* button. The browser will then wait while the packages are securely downloaded and installed on the system. During this time period do **not** hit *stop*, *back* or *reload* in your browser or the packages will not be properly installed.

When the process is complete, a screen displaying a list of all installed packages will be displayed.

The system has now been updated with the latest selected packages available from Guardian Digital.

5.1.4   Run the Installation Agent

The *Installation Agent* is very similiar to the *Update Agent* covered above. Instead of providing updates the *Installation Agent* can perform installations of new packages not originally included in EnGarde Secure Linux, security fixes and bug fixes.

Severity   This will display the severity of the package.

Advisory   This is a link to the text advisory. Clicking on this will open the advisory in a new window.

Available Version This is the latest available version.

Dependencies If all dependencies for this package are met *'resolved'* will be printed here.

6 ENGARDE CONNECTIVITY

So far the only way we spoke of to connect to your EnGarde system was via the GD WebTool utility. To gain remote access you have another secure alternative. We provide SSH connectivity to EnGarde.

Since `telnet` is extremely insecure, it is not provided with EnGarde Secure Professional. SSH uses 1024 bit encryption to protect your connection.

Secure Shell (SSH) is a program for logging into a remote machine, as well as for executing commands on a remote machine. It is intended to replace `rlogin` and `rsh`, and provide secure encrypted communications between two untrusted hosts over an insecure network.

SSH connects and logs into the specified hostname. The user must prove his/her identity to the remote machine using one of several methods depending on the protocol version used. For more information on SSH please visit `www.openssh.com`, the OpenSSH Project home page.

6.1   Connecting from Windows 9x/ME/NT/2000

Windows-based systems only include `telnet` capability. Therefore, we have included a utility to make a secure connection to your EnGarde system from a Windows host. MindTerm is a secure SSH client included on your EnGarde CD-ROM. It can be found in the `x:\dosutils\mindterm` directory. Replace the "x", in the previous statement with the drive letter of your CD-ROM drive. Installation instructions are in the next section.

MindTerm provides you the ability to make an SSH connection to your EnGarde Linux system. You will be on a secure, 1024 bit encrypted connection. MindTerm performs X-Term emulation. You also have SCP capabilities which allows you to copy files securely over an SSH connection. SCP will be fully explained in the *Menus* section.

6.1.1   Installing MindTerm

We have included an installer for Windows based systems to use. You can find the installer in `x:/dosutils/mindterm/setup.exe`. You can type in the command by clicking the *Start* button, then selecting *Run*. You can also click on *My Computer*, select you CD-ROM drive, then the *dosutils* folder, followed by the *mindterm* folder and finally selecting the `setup.exe` file. This will start the MindTerm installer.

Once the installer starts, you will have a few options. You will have to choose the directory you wish to install MindTerm into. The default is `c:\Program Files\mindterm`. We suggest leaving the default. You can then select the installer to create an icon on your desktop for MindTerm and/or an icon in your Start Menu. These are both turned on by default.

Once you have made your selection, select *Install*, which will confirm your selections. If you are satisfied with your settings select *Ok* and MindTerm will start installing. You will see all the MindTerm files scrolling in the window as they are installed. When the installation is done a message box will appear saying: "*MindTerm installation successful!*". You can close this box and now use MindTerm. If you selected the option to install the icon on your desktop you will see it there. If you also had the installer create the Start Menu icon you will find *Start Menu->Programs->MindTerm->MindTerm* and *Readme*. The readme is detailed information about MindTerm and how to use it. We will be covering a general usage of MindTerm in the next section.

NOTE: MindTerm is distributed free. There are other programs for Windows such as
              TeraTerm and Secure-CRT that will also work with EnGarde.

6.1.2   Running MindTerm

MindTerm uses a public/private key cryptography system to connect to EnGarde.
A public key is a key the user is assigned that can be given out to anyone. At the
same time they are also given a private key that no one can have. The public key
is then checked against the private key for authenticity. In the case of EnGarde
Linux the private key is stored on your EnGarde system and MindTerm passes the
public key to EnGarde for authenticity.

You can start up MindTerm by either double clicking on the MindTerm desktop
icon or choosing it from the Start Menu, *Start->Programs->Mindterm->Mindterm*.
After a few moments you will be displayed with the MindTerm screen.

When you started up MindTerm you may have noticed a MS-DOS Prompt window
appear and it may be located behind your MindTerm window. You may minimize
this window but do not close it. The MS-DOS Prompt window will close when
you shutdown MindTerm.

At this point you will need to set up MindTerm so that it knows where to connect

to, who you are and what key to use. First you must have a valid user on the system you are trying to connect to. If you do not have a user, are uncertain of the user name or forgot your password then contact your system administrator. To view and/or modify any of the information mentioned please refer to *Section* 4.4.1 *User Account Administration* on page 77.

You are also required to have a key for the system. The key provides the encrypted information MindTerm requires including your password, to authorize you to connect to the remote host. When your account was created by the system administrator, a key should have been given to you. If you do not have this key please contact your system administrator. To generate a new key refer *Secure Shell Management* on page 84.

To enter this information into MindTerm select *Setting->SSH Connection...*

This will pop up a window labeled "MindTerm - New Server". Here you will need to enter in the information mentioned above. Each field will be described below.

Server In this field you will need to enter in either the IP address or the name of the server you are trying to connect to. In our example above we want to connect to `lockbox.guardiandigital.com`. So `lockbox.guardiandigital.com` was entered in to the server field.

Port This field should be preset to port 22, the default SSH port. We suggest leaving this as is.

Username Here you will need to enter in the user name your system administrator has given you for the server. In our example we are trying to login as user *admin*. This user name will automatically be passed to MindTerm. So you will only need to supply a password when you login. *admin* was entered in to the field.

Cipher In this field you will have a pull-down menu giving you a selection of different cipher methods. A cipher is a method of encrypting plain text information into encrypted information. There are several different methods. By default EnGarde is set to use *3DES*. Check with your system administrator to see if they have changed the cipher.

Authentication Here you will need to select your authentication type. The authentication type is the method that will be used to authenticate you when you log in. By default *RSA* is used. *RSA* uses a public and private key scheme. When your account was created, you should have been given a key to be used with the server. Forms of authentication other than RSA are not supported by EnGarde Secure Professional.

Identity   Here is where you will enter in the path to your key. By default MindTerm
will search in `c:\Windows\Java\mindterm` for keys. It would be ap-
propriate to place your key in this directory when it is given to you by your
system administrator. You can use the "..." button to browse through other
directories on your local machine. A key will generally end with *.key*.

Once all the information has been filled in you, can select the *OK* button to con-
tinue. You will be brought back to the screen you began on.

Once you click the *OK* button MindTerm will attempt to make a connection. If
you have never connected to the server before you will be asked if you want to
add the host to your host key list. Answer *Yes* to this question.

Once the dialog box is removed, if the connection was successful you will be
prompted for your password.

If you do not have the above screen then you most likely received an error. A couple of common errors are:

Unknown Host: You will receive this error if the name or IP address of the host was not found or is not responding. Check what you entered in the *SSH Options* screen above.

Server refused our key You will receive this error if the key you are using does not correspond to the key on the server. This can be caused if the key on the server has changed, you are pointing MindTerm to the wrong key, or your key is invalid. Double check your settings in the *SSH Options*. If you are certain you are passing the correct key, then a new key may have to be generated. Contact your system administrator if this is the case.

At the password prompt displayed above, enter in your password that was assigned to you by your system administrator. If you entered in the password correctly you will now be logged into the system.

At this point you are ready to interact with the system.

Now would probably be a good time to save your settings. Saving your settings allows MindTerm to store the information you entered into the *SSH Connection...* dialog so you don't have to re-enter the data in every time.

To save your settings select *File->Save Settings.*

To exit the system type *exit*. You will be brought back to the SSH Server/Alias: prompt. At this point you can shutdown MindTerm by clicking the 'X' in the corner or from the menu, *File->Exit*.

It is highly recommended that you log out of the server using the *Exit* command before shutting down MindTerm so you are properly logged out.

6.1.3   Secure Copy (SCP)

The Secure Copy (SCP) is a method of copying files over a secured SSH connection. MindTerm supports SCP.

To copy files to and from the server via SCP you will first need to be logged into the system. Read the section above on logging in with MindTerm. You will then have the ability to SCP by selecting *File->SCP File Transfer....*

Selecting the *SCP File Transfer...* option will bring you to the following screen:

This interface works similiar to other FTP clients available for the Windows platform. You can select files be clicking on the filename; multiple files can be selected. Buttons to *create*, *delete*, and *rename* directories. To transfer a file select the arrow facing the machine you want the files transfered to. When doing this you will see a status screen showing the transfer.

Once this status screen reports *Done* the files are completely transfered.

6.1.4   MENUS

The easiest way to learn how MindTerm works and what features it provides is to look through this brief walk-through of all menus in MindTerm. Given within parentheses is the keyboard short-cut for each menu item where one exists.

File Menu

New Terminal (*Ctrl+Shift+N*) This will create a new MindTerm window with the same settings as the first MindTerm window of this session, i.e. all parameters (command-line or applet) given to MindTerm at startup will have effect in each new terminal created.

Clone Terminal (*Ctrl+Shift+O*) This will create a new MindTerm window with the exact same settings as the window it is created from. If the window contains a connected session, the new window will be automatically logged in to the same SSH-server (using the same authentication as was used in the original window). Note that the new window will not have any open tunnels since the window from where it is created have the tunnels opened already (preventing the new window from opening them).

Connect... (*Ctrl+Shift+C*) This launches the Connect dialog. From this dialog you may either select to connect to a host whose settings you have saved or you may create settings for a new host. Note when selecting New Server a new dialog is shown which is identical to the one described in 3.8.1 SSH Connection....

Disconnect (*Ctrl+Shift+D*) This forces the current session to be disconnected. Note that this will cause all tunnels to be closed and the shell to be abandoned without logging out. The preferred way to disconnect is to logout in the shell.

Load Settings... Loads settings from a file (extension .MTP) without connecting to the server.

Save Settings (*Ctrl+Shift+S*) Saves current settings.

Save Settings As... Creates a new settings file and saves current settings to it. Useful for creating a short name for a server, or for having more than one set of settings for a specific server.

Create RSA Identity... Creates an RSA identity to be used with authentication type *rsa* or *rhostsrsa*. Two files are created, one containing the private key (default name *identity*') and one containing only the public key (default name *identity.pub*'). The contents in the file with the extension .pub must be copied to the file *authorized_keys* on the server (typically found in ~/.ssh/). These RSA key-files are identical to the ones used with the Unix version of SSH.

SCP File Transfer... In this dialog you can choose files and/or directories to transfer to or from the SSH-server. Local file(s)/dir(s) is a space-separated list of files and/or directories (if a name contains a space enclose it in quotes like: *a file with spaces*). Normal regexp's can't be used for local files/dirs, however names can be given with ONE wild-card ('*') in it (e.g. *.foo or foo*bar). If absolute path-names are not given the current directory is assumed (defaults to MindTerm's home-directory). If the first file/directory given contains an absolute path-name this directory is used as current-directory for the rest of the list (e.g. the list /tmp/foo* *.bar will expand to all files starting with FOO or ending with .BAR in the directory /tmp'). Remote files(s)/dir(s) are given EXACTLY as they would be with the standard Unix scp-client (i.e. regexps can be used). The directory assumed on the

remote side is the user's home-directory (i.e. just like with the standard unix scp-client).

To change direction of the copy-operation press the *Change Direction* button (the direction is indicated with the strings (source) and (destination) after the respective side.

If directories are to be traversed enable *Recursive copy*. To make the copy-operation use as little bandwidth/CPU as possible set it to be *Low priority*. Press *Start Copy* to start the copy operation. This will launch a small window with progress and statistics of the copy operation. A copy-operation can be canceled at any time by pressing the *Cancel* button in this window.

Capture To File... Captures terminal-output to a file. Capture starts immediately when the file has been selected and ends when this menu item is selected again. Note that while capturing is active this is indicated by the menu item being selected.

Send ASCII File... This will send the contents of the selected file to the terminal as input (i.e. would be the same as if the contents were typed from the keyboard)

Close (*Ctrl+Shift+E*) Closes this window. Note that when closing a window without logging out you are aborting the SSH-connection abnormally, i.e. it is advisable to logout in the shell before closing/exiting MindTerm.

Exit (*Ctrl+Shift+X*) Closes all windows and exits MindTerm. Note that when closing windows without logging out you are aborting the SSH-connection abnormally, i.e. it is advisable to logout in the shell before closing/exiting MindTerm.

Edit Note, the system clip-board is not available to applets by default. In this case a local (to MindTerm) clip-board is used. Also note that in some implementations of the Java runtime the clip-board does not work with the system clip-board.

Copy (*Ctrl+Ins*) Copies selected text to clipboard. Selection is done by clicking and holding down left mouse-button while dragging the mouse over the area to select.

Paste (*Shift+Ins*) Pastes the contents of the clipboard to the terminal as input (i.e. would be the same as if typed from keyboard)
Copy & Paste Does a copy followed by a paste.

Select All (*Ctrl+Shift+A*) Selects all content in scroll-back buffer and in terminal. Note, this operation is very time-consuming right now.

Find... (*Ctrl+Shift+F*) Shows Find dialog from which the scroll-back buffer and terminal contents can be searched for words. The search can be done case sensitive or case insensitive. Each word found is highlighted. The bell is sounded when no more matches is found.

Clear Screen Clears screen and sets cursor position to upper left corner.

Clear Scrollback Clears contents of scroll-back buffer.

VT Reset Resets terminal-settings to default (e.g. clears line-draw graphics mode which might be mistakenly set by displaying a binary file).

Settings

SSH Connection... (*Ctrl+Shift+H*) In this dialog you can set all SSH parameters. To view all options click the button *More options...*. When connected you can set the parameters for the current session. Note that some changes wont take effect until the next time you connect to this server. When not connected a new session is created if one is not found with the name of the server. In this case it is the same dialog that is shown when selecting *New Server...* from the Connection dialog .

The parameters set in this dialog are (names as given in paragraph 5.):

```
server   Name (ip-address) of SSH-server port

Port     which SSH-server listens on username

User     name to login as on SSH-server

cipher   Name of block-cipher to use, or if none is
         selected no encryption (note, no encryption is
         normally not supported by the SSH-server)

authtyp  Method of authentication, or if custom...  is
         selected a comma- separated list of methods to
         try in order given
```

```
x11fwd   Selects whether to allow X11-connections to be
         forwarded or not

display  The local X11 display to forward X11 connections to

mtu      Maximum packet size to use alive Keep

alive    interval in seconds to use

portftp  Enables port-commands to be used with
         FTP-tunnels, don't enable this if you are not
         sure what you are doing

realsrv  Real ip-address of SSH server if it is behind
         address translation (used when portftp is enabled)

localhst Address to listen on for local tunnels

idhost   Sets whether to verify identity of the
         SSH-server using its host-key through matching
         with saved value in the file known_hosts

forcpty  Force allocation of PTY, e.g.  necessary to
         enable when executing a single command on the
         SSH-serverthat requires a non-dumb terminal

prvport  Used to force the local outgoing port
         of the connection to the SSH-server to use
         a so called privileged port (i.e.  < 1024)

remfwd   Enables hosts other than the one running
         MindTerm to connect through SSH-tunnels
```

Terminal... (*Ctrl+Shift+T*) In this dialog you can set the basic terminal parameters, such as terminal type, size, font and colors. The initial window position can optionally also be set. It is given as a string with the syntax <+/-><x-position><+/-><y-position> a negative sign means it's relative to the right or bottom. A value of zero means aligned to the border (i.e. left, right, top, bottom) e.g. +0-0 means aligned to bottom right corner.

The parameters set in this dialog are (names as given in paragraph 5.):

```
te       Terminal type
```

```
gm        Terminal geometry, number of lines,
          columns and optionally initial position

fg        Foreground color, name or when custom rgb
          is selected an rgb-value

bg        Foreground color, name or when custom rgb
          is selected an rgb-value

cc        Cursor color, name or when custom rgb is
          selected an rgb-value
```

Terminal Misc... (*Ctrl+Shift+M*) This dialog contains some extra settings for the terminal.

The parameters set in this dialog are (names as given in paragraph 5.):

```
sl        Number of lines to save in scroll-back buffer

sb        Position of scrollbar, or disable scrollbar

sd        String containing delimiter characters that
          are used when click-selecting words, i.e.
          which characters functions as word-delimiters

bs        Indicates whether backspace or delete should
          be sent when backspace-key is pressed

de        Indicates whether backspace or delete should

be        sent when delete-key is pressed
```

Local Command-Shell Starts the local command-shell from which one can view and set all parameters of MindTerm. The command-shell is really only useful if you don't have menus (e.g. when running without a GUI) but for completeness it is available here. Note, the command-shell is only available if enabled with command-line option *–c* or applet-parameter *cmdsh*.

Auto Save Settings Enables/disables automatic saving of settings, when disabled you must explicitly save settings to file when needed. When enabled settings are saved whenever you disconnect from a server or when you exit

MindTerm. Note that when both auto-save and auto-load is enabled (which is default), settings-files are created automatically and the user never have to worry about saving/loading them.

Auto Load Settings Enables/disables automatic loading of settings. When disabled you must explicitly load settings from file if you need to. When enabled, MindTerm tries to load a settings-file with the same name as what you give at the *SSH Server*: prompt or in the (*Settings -> SSH Connection...*) dialog. These files are located in the MindTerm home-directory. Thus the server you give at the prompt does not necessarily have to be the name of the server, it is mainly the name of the settings-file to load. Normally the user does not have to worry about the settings-files since it is handled automatically. Though to create short-names for servers and to create multiple settings-files for a single server you have to explicitly create settings-files.

Current Connections... This dialog lists the currently open connections through the tunnels you have set up. Note that it doesn't list the tunnels themselves, only active connections through them. You can close a tunnel by selecting it and clicking close.

6.2   Connecting from Unix

The first thing you will need to connect to your EnGarde system is an SSH client. For Unix there is OpenSSH. You can download OpenSSH from `http://www.guardiandigital.com/tools`. You will also find OpenSSL, as you will need this too. If you wish to download OpenSSL you can find it at `http://www.guardiandigital.com/tools`. A version of OpenSSL and OpenSSH are included on the EnGarde CD-ROM.

If you are using Windows, use the included MindBright MindTerm software. You can find it on the EnGarde CD-ROM under the *dosutils* directory. Instructions on installation and usage can be found in the previous section.

6.2.1   Using OpenSSH

The first thing you will have to do is create a user. This is either done by logging in as root at the console and running *adduser* or adding a user from the GD WebTool utility.

If you use the GD WebTool utility to create the user read *Section* 4.4.1 *User Account Administration* on page 77 on how to accomplish this.

If you decide to create the user from the console use the following steps:

As the root user run *adduser* by typing *adduser* at the prompt. *adduser* will prompt you for a user name. Enter the user name you wish to give this user.

Once this is done you will be back at the prompt. You now need to give this user a password for them to use to access their account. Type *passwd username*. In place of *username* will be the user name you assigned to the user. This will prompt you for a password and then prompt you again for the password to confirm it.

Once that is done install OpenSSL and OpenSSH on your client machine.

NOTE: You must be root during the installation of OpenSSL and OpenSSH.

On distributions using RPM:

```
$ rpm -Uhv openssl-0.9.4_i386.rpm
$ rpm -Uhv openssh-1.2.3_i386.rpm
```

In Debian (or any distribution using DPKG):

```
$ dpkg -i openssl-0.9.4.dpkg
$ dpkg -i openssh-1.2.3.dpkg
```

And from tar files:

```
$ tar zxvf openssl-0.9.4.tgz
$ tar zxvf openssh-1.2.3.tgz
$ cd openssl-0.9.4
$ ./configure
$ make
$ make install
$ cd images/openssh-1.2.3
$ ./configure
$ make
$ make install
```

You now must create a key for yourself. You can create a key with OpenSSH by typing:

```
$ ssh-keygen
Generating RSA keys: ......oooooooO.................oooooooO
Key generation complete.
Enter file in which to save the key (/home/nick/.ssh/identity):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
```

It will prompt you for a filename to save the key in. The default `identity.pub` will be fine. It will then prompt you for a new passphrase. After entering your passphrase twice, your public key will then be generated.

Once you have your key e-mail it to your system administrator and they will insert it in to the system properly. Read *Section* 4.4.3 *Secure Shell Management* on page 84 for more information. Once this has been completed you will be able to successfully SSH in to the system.

For more information on SSH and using SSH please read the SSH FAQ which can be found at:

```
http://www.linuxsecurity.com/docs
```

7 VIRTUAL PRIVATE NETWORKING

EnGarde Secure Professional and the accompanying Workgroup Suite implement Virtual Private Networking (VPN) using the PPTP protocol. The Point-to-Point Tunneling Protocol is a network protocol that enables remote office workers to connect to their local corporate network behind their EnGarde Secure Professional gateway server, protecting their communications through a secure private data channel. EnGarde Secure Professional employs sophisticated encryption technology to ensure that data transmitted from the remote workstation to the EnGarde gateway cannot be intercepted and remains secure during its transmission.

Using PPTP on EnGarde Secure Professional, remote office workers can connect to their internal hosts to access network resources such as file and e-mail services.

EnGarde Secure Professional implements a standards-compliant PPTP server implementation that supports Windows 98, Windows NT, and Windows 2000 clients. While support for the PPTP protocol is included in Windows NT and Windows 2000, it must be downloaded and installed for use with Windows 98.

For an example of how PPTP might be used to provide VPN services in your organization, refer to Figure 5 on page 199. Details of the PPTP protocol itself and additional information are available by searching microsoft.com for "Understanding PPTP" document dated January 1997.

NOTE: Virtual Private Networking is only available with the purchase of the EnGarde Workgroup Suite.

The following text description and corresponding image depict a typical PPTP session of how a remote user might connect to their corporate network.

(A) The Windows PC client will make a PPTP connection using an existing connection to the Internet. PPTP will encrypt the data before sending it out over the Internet to the EnGarde Secure Professional server at the other end.

(B) The EnGarde server is the destination for the PPP packets containing the encrypted PPTP information within.

(C) When the EnGarde server receives these packets it will decrypt the information and distribute it to its destination within the local network. An additional IP address will be assigned by the EnGarde sever on the internal network to the remote Windows workstation.

Figure 5: PPTP general overview

(D) At this point you can access all of your internal network resources as if you were locally connected to the network. You have access to your e-mail account, ability to send e-mail from the network, access the internal only Intranet among many other tasks.

7.1   Configuring EnGarde for PPTP File and Print Sharing

To allow outside users to access internal resource shares on your EnGarde Secure Professional server through a PPTP connection you must have both *Local Master?* and *Allow Domain Logins?* set to **Yes** in the *System Management->Windows File Sharing->Global Configuration*.

NOTE: For a full description of the WebTool PPTP interface refer to *Section 4.4.7* on page 111.

Next you must make certain that in *Security->PPTP Setup->General Security* you have the *Local WINS Server* set to the IP address of your EnGarde Secure Professional machine. In our example we are using `192.168.1.82` as our EnGarde server.

While in the *PPTP Setup* section of the WebTool, make certain you have a user account so that the remote user has access to login.

Finally, you must restart PPTP for the new changes to take effect.

Select *System Status Monitor* from the main WebTool menu. Then select *Services Monitor* from the *System Status Monitor* menu. This will display a list of the available services. Toggle the status of the service by clicking on *Enable* or *Disable*.

NOTE: For detailed information concerning use of the *System Status Monitor* refer to *Section 4.5* on page 134.

The PPTP server has now been configured and restarted. You are now ready to configure your Windows clients.

7.2   Connecting From Windows 98

You can find many of the necessary system updates using Microsoft's Windows Update technology and the Internet Explorer Web browser from the Windows 98 client machine.

Listed below are the required packages for PPTP to successfully operate as well as a list of recommended packages. They can be obtained by accessing http://windowsupdate.microsoft.com using Internet Explorer only.

The recommended packages are not necessary, but on some older versions of Windows 98 may be required, and will also improve performance.

Windows Update Required Components:

- 128-bit Encryption Pack

- Internet Explorer 5.5 or greater

- Root Certificate update

Microsoft frequently also issues "Critical Update Packs" through the Windows Update facility. It is recommended that all critical updates are also installed, as these often fix security vulnerabilities that may prevent system compromise.

Once these components have been successfully installed, it is necessary to update Microsoft Dial-Up Networking (DUN) to at least version 1.4 by reading the following Microsoft document and following the instructions within:

`http://support.microsoft.com/support/kb/articles/Q285/1/89.ASP`

Windows 98 Setup

Once the updates have been completed you are ready to set up the connection to your EnGarde Secure Professional PPTP server.

To set up PPTP in Windows 98 start by clicking on *My Computer* on your desktop.

The PPTP protocol in Windows 98 uses the *Dial-Up Networking* interface. Create a new "connection" in *Dial-Up Networking* by clicking on the *Dial-Up Networking* icon.

Select the *Make New Connection* icon to start the connection wizard application.

There are two options on this first screen. The first is labeled "a name for the computer". This is just a label that will be associated with this new configuration. After the PPTP connection configuration is created it will be displayed as an icon with the label you give it below it. The *My Connection* default name can be changed to something more descriptive such as *Corporate Network*.

The second option here is a pull-down menu box. Since we are setting up a Virtual Private Network (VPN), you will want to select *Microsoft VPN Adapter*.

When all your changes are set hit *Next* to proceed.

The next step is to enter in the IP address of the EnGarde Secure Professional machine.

This is the last step of the creation process and the wizard will ask you to confirm everything. If everything is set up properly, click the *Finish* button and the process is complete.

Once the connection is defined it will be added into the *Dial-Up Networking* folder. You will see it listed with the name you gave it below.

NOTE: By dragging the *My Connection* icon to the desktop a link will be created to make it easier to access.

Before attempting to establish a connection a couple settings must be confirmed and possibly changed first. Go into the properties of the new PPTP configuration you just created by right-clicking on it and selecting *Properties*.

The following screen will appear. Make sure your screen has the same options configured as this one. Disregard the default TCP/IP settings found by clicking on the *TCP/IP Settings* button.

Once the necessary changes are made click *OK*.

You are now ready to attempt to establish a connection. Double left-click on the configuration you created, *My Connection* in the example used above. You will see the *Connect To* dialog box appear. Enter in the user name and password you set up on the EnGarde Secure Professional machine into these entry boxes.

Once the information has been entered click *Connect* to establish a connection with the EnGarde Secure Professional server.

NOTE: It is recommended you reboot your Windows system before attempting to connect.

As Windows attempts to make the connection you will see the *Connect To* dialog box replaced with a smaller dialog box displaying the results of the connection. If the connection is successful you will see what appears to be an icon of two computers connected together in your task bar. Each "computer" will light up green when data is sent and received over this connection.

By double-clicking on the computer icon in your task bar, a status dialog box showing information about your PPTP connection will be displayed. You can get detailed information concerning the protocols by clicking the *Details > >* button, disconnect from the network with the *Disconnect* button or hide the dialog box by clicking *OK*.

NOTE: Do not confuse this with a dial-up connection using a modem. This is connecting to another network over an existing connection.

You can now access the resources on the network you connected to via the *Network Neighborhood*.

7.3   Connecting From Windows NT 3.5

To configure PPTP to work in Windows NT 3.5 you will need to first install the PPTP drivers from the network menu in the Control Panel if they aren't already installed. Use the *Add/Remove Programs* section in the *Control Panel* for this or contact your system administrator.

NOTE: The Windows NT 3.5 CD will be required to install the PPTP drivers.

Once the PPTP drivers are installed *Service Pack 6a* is required to be downloaded and installed.

After Service Pack 6a has been installed and the system is rebooted you are ready to start configuring your Windows NT 3.5 machine to connect to your EnGarde Secure Professional with PPTP.

Windows NT 3.5 uses the *Dial-Up Networking* interface to control PPTP. A new Dial-up configuration will be created for your PPTP connection. To create this configuration click on *My Computer*. From the *My Computer* window, select *Dial-up Networking*.

Dial-up Networking will start up with your dial-up configuration(s). If no other configurations were made previously, then the fields will be empty, as in the example below.

Click the *New* button to start the creation of a new dial-up configuration.

The first step is to give this dial-up configuration a name. For the example *My-Connection* was used. Anything descriptive can be used, blank spaces are not allowed.

Click *Next* to continue.

Next we need to tell Dial-up Networking how we are going to go about establishing our connection. The PPTP service will use an existing connection to the Internet as a passageway to the remote network. The option in Windows NT is called "I am calling the Internet".

Click the *I am calling the Internet* check-box and click *Next* to continue.

Next the information concerning what computer you want to connect to must be entered. A PPTP connection to an EnGarde gateway does not require a phone number here, but rather an IP address. Enter in the IP address of the EnGarde gateway into the *Phone number:* entry-box.

At this point your new PPTP configuration is complete.

Hit *Finish* to write the configuration.

You will now be returned to the Dial-Up Networking section with the option to *Dial* with your new configuration.

NOTE: The Dial button will not physically dial another computer but makes a connection to another network via a currently established Internet connection.

Click the *Dial* button to connect to your EnGarde Secure Professional server.

The *Connect To* dialog will appear. Enter in your user name and password you selected when creating the user account on your EnGarde machine.

Click the *OK* button.

Windows will attempt to establish a connection to your EnGarde machine using PPTP. When a connection is established an icon will appear on your task bar and a "bubble" containing our connection information will appear for a few moments.

You can bring up a status screen and other options be double-clicking on the icon in the task-bar.

You are now ready to access other shares and other network resources.

7.4   Connecting From Windows 2000

Windows 2000 was designed with the PPTP protocol built-in and no updates or patches are required specifically for PPTP.

NOTE: It is always recommended you have the latest service packs released by Microsoft installed to reduce possible problems.

To setup a PPTP connection to your EnGarde Secure Professional PPTP server, start by clicking the *Start* button. From there select *Settings->Network and Dial-up Connections->Make New Connection*.

This will start the *Network Connection Wizard*. Click *Next* to start the PPTP configuration process.

The first configuration option here is to choose which type of connection you will be making. We want to setup a VPN (Virtual Private Networking) connection. So select the *Connect to a private network through the Internet* option.

Click the *Next* button to continue.

If you need to connect to an ISP or use a dial-up connection of some type to get on the Internet, the PPTP configuration can be set up to automatically dial your Internet connection for you before trying to establish a connection to the PPTP server. To configure it to do this choose your connection from the list-box, otherwise choose the first option.

Click the *Next* button to continue.

This next dialog box requires only that you enter in the IP address of the EnGarde PPTP server to make your connection.

Click the *Next* button to continue.

Finally assign a name to label this connection. You can also choose to have it create a link to this connection on your desktop.

Click the *Finish* button to create this connection.

After creating the new connection Windows 2000 will automatically display a dialog box to establish the connection. We do not want this done just yet as a couple other settings need to be confirmed.

The following icon is created on your desktop if you chose to have the *Connection Wizard* create it. Right-click on the icon and select *Properties*.

If you chose not to create the icon select *Start->Settings->Network and Dial-up Connections->"Your new connection"*, right-click on it and select *Properties* from there. The *Properties* dialog will be displayed.

In this new dialog select the *Networking* tab. Make certain your properties have the same configuration as the one below has. Hit *OK* to accept the changes you may have made.

We are now ready to attempt to establish a connection. Double left-click on the icon. The connection dialog box will be displayed prompting you for a user name and password. Use the user name and password you configured through the WebTool previously.

Once this information has been entered into the entry boxes select *Connect* to make the connection.

If the connection is successful an icon of what looks like two computers connected together will appear on your task-bar. You can click on this icon to get statistics about the connection and to terminate the connection.

You will also notice the icon on your desktop will change, if you selected to create a desktop icon. The monitors on the two computers in the icon will turn from gray to blue informing you that a connection is established with that PPTP configuration.

You are now connected to your inside network and have access to all the resources. Use the *Network Neighborhood* to access files and printers.

8   SECURE E-MAIL

EnGarde Secure Professional provides two methods of retrieving your e-mail remotely, secure IMAP and secure POP3. Both protocols have been secured using SSL and both require clients that support SSL secured IMAP and secured POP3.

Securing IMAP and POP3 greatly increases the security and privacy of personal e-mail. For this reason IMAP and POP3 are only available in a secure form and therefore the standard, insecure form of IMAP and POP3 are not available with EnGarde.

Using a secure form of these protocols requires a client that can support them. We will discuss how to configure both Netscape Mail for secure IMAP and Microsoft Outlook for secure IMAP and secure POP3.

8.1   Configuring Netscape Mail for Secure IMAP

The Netscape Communicator package includes Netscape Mail. Netscape Mail is capable of both IMAP and POP3 but only supports IMAP in secure mode. Below is a set of instructions for configuring your Netscape Mail for secure IMAP.

NOTE: You must allow users to access their mail from their machine by adding in their IP address in the *System Access Control Section 4.6.5* on page 144.

To access the Netscape Mail you will first need to start Netscape. Once Netscape is loaded you can launch the Mail by either selecting *Communicator->Messages* or by clicking the *mail* icon in the lower corner of the browser window.

At this point the Netscape Mail window will appear. Now pull-down the *Edit* menu and select *Preferences* from there.

After selecting *Preferences* the *Preferences* window will be displayed. From here you will want to expand the *Mail & Newsgroups* section by click on the '+' found in the box. You will then have a new group of options. We will start by configuring our user name, e-mail address, etc. Click the *Identity* option from the menu tree on the left.

Once the window appears fill in the appropriate information. When you are done entering everything select *Mail Servers* from the menu tree on the left. This will bring up the options for your incoming and outgoing e-mail servers.

We will start be creating a new server for the incoming mail. First delete the
default server Netscape includes by clicking on it and selecting the *Delete* button.
Then click the *Add* button.

You will be presented with the following dialog:

In the *Server Name* field you will need to enter in the name of the mail server given to you by your system administrator. In the example above we used `lockbox.guardiandigital.com`.

Next we need to select the *Server Type*. Netscape Mail only supports secure IMAP so select *IMAP Server* here.

Finally in the *User Name* field enter the user name you were assigned to by your system administrator.

Next click the *IMAP* tab at the top of the dialog. You will be presented with a number of IMAP options.

Here you will want to make sure all the check-boxes are turned off except for the User secure connection (SSL) option. Your screen should match the number above.

After closing the *Mail Server Properties* dialog you will see your mail server in the window labeled *Incoming Mail Servers*. Finally you will have to enter in the server name for your outgoing e-mail. Enter in the outgoing server name given to you by your system administrator in the *Outgoing mail (SMTP) server* field and enter your user name in the *Outgoing mail server user name* field.

Once you have completed entering in the information click the *OK* button. The Preferences dialog will close and you will see the server name appear in your mail listing, where your Inbox is located.

You are now ready to receive mail from your EnGarde Linux system with Netscape Mail using secure IMAP.

NOTE: You must allow users to access their mail from their machine by adding in their IP address in the *System Access Control Section 4.6.5 on page 144.*

8.2 Configuring Outlook for Secure IMAP and POP3

Microsoft Outlook 2000 is capable of both IMAP and POP3 and supports both protocols in secure mode. Below is a set of instructions for configuring Outlook 2000 for secure IMAP and POP3.

NOTE: Outlook 2000 is required. Previous version of Outlook do not support these features and will not work.

NOTE: You must allow users to access their mail from their machine by adding in their IP address in the *System Access Control Section 4.6.5 on page 144.*

Begin by starting up Outlook. Once Outlook is loaded you can create a new e-mail profile by selecting the *Tools* menu and from there select *Options*.

NOTE: If this is the first time you are using Outlook it will automatically start in the Internet Connection Wizard section to create an e-mail profile. If this is the case skip down in this section to the Internet Connection Wizard and start from there.

At this point you will be presented with the *Options* screen. From here select the *Mail Delivery* tab and click the *Accounts* button from within there.

You will now see the *Internet Accounts* dialog. Our objective is to create a new e-mail profile first with basic information. Then edit the profile to allow for secure POP3 or IMAP. So here we want to add the profile, so click the *Add* button.

You will now be prompted with a small "pull-down" type menu. You have two options in here *Mail and Directory Service*. Since we are creating a new e-mail profile select the *Mail* option.

Now you will see the *Internet Connection Wizard* start. The *Internet Connection Wizard* will go through a step-by-step process to create the basic account. Once the basic account is created we will have to edit the account to accept secure e-mail transfers.

The first step in the *Internet Connection Wizard* is to enter your full name. This is the name that will be automatically displayed when someone receives e-mail from you.

Once you have entered your name in click the *Next* button to continue.

Now you will be prompted for your e-mail address. This has most likely been assigned to you by your system administrator.

Once you have entered in your e-mail address click the *Next* button to continue.

You will now be presented with a few options. You first have the choice of using POP3 or IMAP for your connection. Select this according to what your system administrator recommends you use. For the remainder of this example we will be using POP3.

You now have to enter the mail server you will be contacting. In our example below our incoming mail server is the same as our outgoing server. In many situations `smtp.servername.com` and `mail.servername.com` are used for outgoing and incoming mail servers.

Once you have entered in the proper mail server addresses and selected the POP3 or IMAP protocol click the *Next* button to continue.

Now you will need to enter in some account information. First enter in your account user name assigned to you by your system administrator followed by the password. You can select the *Remember password* option if you wish for Outlook to remember the password for future sessions.

You will also notice a check-box for *Secure Password Authentication (SPA)*. This feature isn't used with EnGarde so leave it unchecked.

Once you have correctly entered in all the required information click the *Next* button to continue.

Now you will need to select which method you use to connect to the Internet. Select the appropriate option and then click the *Next* button to continue.

You will now see a confirmation screen informing you the profile has been created. Click the *Finish* button to continue.

You will now be returned to the *Internet Accounts* dialog and will notice the profile you created listed in the window in the *Mail* tab. At this point we have to setup the profile to work with a secure server. Select the *Properties* button on the right.

Here you will see you have four tags, *General*, *Servers*, *Connection* and *Advanced*.
Select the *Advanced* tag to continue.

You will now see a number of options in this screen. We are only concerned with the options displayed below the *Server Port Numbers* section. You will want to select the box below *Incoming mail (POP3)*, this will say *(IMAP)* if you selected IMAP as your server. Once you click the box you will see *995* appear in the text field, or *993* if you selected IMAP instead of POP3 earlier. At this point you can click the *OK* button to finish.

Your Outlook mail client is now configured to receive secure e-mail via POP3 and IMAP.

NOTE: You must allow users to access their mail from their machine by adding in their IP address in the *System Access Control Section 4.6.5* on page 144.

design by t.lum

9 THE LINUX INTRUSION DETECTION SYSTEM (LIDS)

9.1 Introduction to LIDS

With the rapid pace of development and open source nature of Linux, programs are often evaluated for security vulnerabilities. Between the time the known security vulnerabilities are found, additional protection is available to provide an extra layer of security, until the system can be updated.

Since Linux is an art of the open source community, security holes may be found more easily but can also be patched just as quickly and easily. But when the hole is disclosed to the public, and the administrator is unable to patch the hole, it could potentially compromise your system. With the typical Linux systems, a cracker has absolute control if superuser access is gained. With the added protection of LIDS, this and many other potential problems can be reduced.

LIDS provides the ability to control all access to files, processes, binaries, memory, raw devices, drives, etc. One of the main features of LIDS is protection from the superuser, known on a Linux system as the root user.

NOTE: LIDS requires advanced administration skills to manage properly and therefore should not be modified by inexperienced users. Managing EnGarde Se-

cure Linux through the WebTool will not require users to perform and LIDS administration.

The root user has control over every single aspect of the system. They can mount and unmount drives, delete and create files, remove users, access the database, edit the Web page, shutdown the system, etc. So you can see the possible security hazard here. If someone managed to gain root access, the entire system could be put into the crackers control. Here is a number of security enhancements LIDS has to protect the system from this threat.

- Every single file can be protected. Giving each file its own set of read, write, or append rules that even the root user must obey. For example, if you set your log files to append only, no one could go in and delete any trace of themselves on the system. You can set the login binary as read-only and it can not be replaced. Even if there was a possible way to overwrite the file LIDS would know it's not the same file because it indexes the files by their inodes, not their file names.

- Files can also be completely hidden from view and only be accessible by specific programs. For example, if you want to protect your Apache SSL server key from everyone including root, you can hide the file so to every user, including root, it doesn't exist, but at the same time it allows Apache to have full access to the file so it can get the information it needs from it.

- LIDS can also protect processes from being killed by the root user. This could be used to protect your database server, your Web server, your mail server, etc. from being taken off-line by an intruder.

- You can have full control of the Linux kernel "capabilities". The current Linux capabilities control what a process can and can't do. Changing these capabilities gives you more control over your system. By setting the capabilities to your needs you can prevent all users from rebooting the system, mounting and unmounting disks, changing network settings, `/dev` control, ownership control, loading and unloading of kernel modules, and many others.

- Root has the ability to turn LIDS off locally for just the current session or globally. This can be configured so it can only be done locally, and/or remotely. It also requires a password which is protected by Ripe MD-160 encryption.

- A built in port scanner allows you to disable promiscuous mode and still detect port scans.

- All attempts on the system are logged and if any user tried to break one of the LIDS rules, an e-mail is immediately sent to a predefined e-mail address. (A cell phone or a pager can be configured to be alerted when this happens also so you know when someone is making an attempt on your system.)

Some minor drawbacks to this increased method of security is it could hinder the use of certain programs by denying them access to needed files if configured incorrectly. It also makes it more difficult to administer the system from the console but the included GD WebTool includes enhancements that integrate will with LIDS.

9.2   Using LIDS

LIDS, be default, is always running on your EnGarde system. If you will be doing your administration via the GD WebTool you can skip this section, but it's suggested reading anyway.

Minimal maintenance is required to keep LIDS running. Management of LIDS on servers that are co-located with Guardian Digital is included with your support contract.

You may sometimes need to change the configuration or add new packages requiring you to disable LIDS. The GD WebTool will automatically enable and disable LIDS while you administer the system. For administration from a shell, a program called lidsadm is used to interface with LIDS.

First you have to disable LIDS. After logging in as root type:

```
/sbin/lidsadm -S -- -LIDS
```

This will prompt you for your password. After entering your password LIDS is disabled for the current session you are in. This method will still apply all the LIDS resource settings and rules to every other user on the system while you administer the system. Optionally, issuing:

```
/sbin/lidsadm -S -- -LIDS_GLOBAL
```

will disable LIDS globally. While in this mode no LIDS rules will be applied to any user or resource. Use this with caution. Once you have LIDS turned off you may configure your capabilities, file permissions, resource permissions, etc. If you changed the LIDS configuration while LIDS was turned off you will need to reload the configuration file into LIDS. Before turning LIDS on enter this:

```
/sbin/lidsadm -S -- +RELOAD_CONF
```

This will make sure you have the latest configuration loaded into LIDS. It is suggested you run this command every time you make a change to the LIDS configuration. To turn LIDS protection back on after administration simply issue:

```
/sbin/lidsadm -S -- +LIDS
```

or to enable it globally:

```
/sbin/lidsadm -S -- +LIDS_GLOBAL
```

Your system is now protected again by LIDS. When enabling, disabling and reloading the configuration information with lidsadm you will be prompted for a password every time. You will see the following message:

```
SWITCH

WARNING: Only system administrators should enable/disable
LIDS. Disabling LIDS can open your Lockbox to possible at-
tacks.  Make sure you read the LIDS section in your in-
cluded manual before manually changing options in LIDS.
Incorrect configurations can have drastic effects.

enter password:
```

At this point you can enter in your password.

9.2.1   Using the lidsadm Utility

The lidsadm utility is a small program you will use to administer your LIDS configuration. It stores all configuration information in /etc/lids/lids.conf. If you are using the GD WebTool for administering LIDS you do not need to use lidsadm.

Some basic lidsadm options are as follows:

```
/sbin/lidsadm -A Add a new entry

/sbin/lidsadm -D Delete an entry

/sbin/lidsadm -Z Delete all entries

/sbin/lidsadm -U Update all entries

/sbin/lidsadm -L List current entries, requires LIDS to be turned off

/sbin/lidsadm -P Creates a new password.  It will store the password
          in Ripe MD-160 encryption

/sbin/lidsadm -S Switch LIDS on/off and capabilities
```

```
/sbin/lidsadm -r View current status of LIDS
```

```
/sbin/lidsadm -h Help
```

The next section will contain more detailed information about the lidsadm options

9.2.2   Adding an Entry

Using this option allows you to add a new item to the LIDS config. You have the options to add a single file with an attribute, give a file permission to override another files permissions, and change the capabilities of a file.

```
lidsadm -A [-s subject] -o object [-t] -j TARGET
```

To protect a file enter the filename and path using the *-o* flag, followed by the attribute, READ, WRITE, IGNORE, DENY, or APPEND under the *-j* attribute. If your object is a capability setting you need to use the *-t* flag to tell lidsadm it's a special option. *-s* is used to point the object to a subject. In the case of capabilities you, are pointing a capability to the subject or giving the subject the capability. Same idea with file protections. If you deny access to a file but want the subject to use it, you point to the denied file(*object*) to the file to give access to(*subject*) then tell it what kind of access to give it *-j*. Here's an example of protecting a file:

```
lidsadm -A -o /path/to/protected_file -j DENY
```

Now to give a binary full access to the file that was denied to everyone else:

```
lidsadm -A -s /path/to/binary \
           -o /path/to/protected_file -j WRITE
```

We also want to give the binary the capability to chown, which has been disabled earlier by LIDS:

```
lidsadm -A -s /path/to/binary \
        -t -o CAP_CHOWN -j INHERIT
```

When changing a files capabilities we use INHERIT or NO_INHERIT instead of the READ...APPEND commands. Using INHERIT gives the file access to the capability while the NO_INHERIT turns off the files abilities to use the given capability. In a later section capabilities are explained in more detail. In the next session an example of a package being protected is given.

NOTE: Don't forget to do a *lidsadm -S – +RELOAD_CONF* after changes were made
             so they take effect when you reload LIDS.

9.2.3   Deleting an Entry

Deleting an entry is an extremely simple task and there is no need to go into great
detail. If there is a file you no longer want to be protected or wish to change
protection on, you need to delete the entry from the LIDS config. Simply issue
the following command to accomplish this task:

```
lidsadm -D [-s file] [-o file]
```

and the file will be removed from the configuration. You can now enter new
attributes for the file, if you like.

9.2.4   Deleting and Updating All Entries

Lidsadm gives you the ability to delete and update all the file entries in your con-
figuration. Issuing:

```
lidsadm -Z
```

will delete every entry in your LIDS configuration and you will be starting with a
clean configuration file. The original configuration shipped on your box is stored
in */usr/bin/lids_default_config/* and can be executed to revert LIDS back to it's
original configuration.

Updating all the file entries works a little differently. The configuration files are
linked to LIDS by their inode number, not their filename. If a file gets deleted
and replaced later it may not be protected by lids because of the inode change. By
issuing:

```
lidsadm -U
```

lidsadm will go through your configuration and check every file making changes
as necessary. This should be ran if you upgrade a package too since it's more than
likely one or more of the files will be overwritten and the inode will change.

9.2.5   Password Creation

LIDS uses a user defined password it stores in encrypted form(Ripe MD-160), in
`/etc/lids/lids.pw`. To create a new password simply type:

```
lidsadm -P
```

It will prompt you twice for your new password and then change the password.
This will obviously only work if LIDS is turned off. Once you have done this
every time you need to reload the configuration and turn LIDS on or off you will
have to enter your password in plaintext.

9.2.6   Viewing LIDS Status

You can use:

```
lidsadm -r
```

to view the current running status of LIDS. This can be useful for writing scripts
that need to know if LIDS is turned on or not.

9.2.7   Viewing the Current LIDS Configuration

You can use the:

```
lidsadm -L
```

option to view a list of all the files and their attributes in the configuration. You
must have LIDS disabled to run this command since it requires access to the
`/etc/lids/lids.conf` file.

9.3   Protecting Your Files

EnGarde Secure Professional comes with a default configuration for protecting your files based on your configuration options and installed packages. If packages are removed, or added LIDS will have to be updated. Most of this can be easily accomplished using the GD WebTool application.

If you wish to do administration of LIDS from the console you will need to use the lidsadm program. Using the commands described in the previous section we will remove, add and update files on your EnGarde system Before any administration can be done you must first turn off LIDS. Turn LIDS off only on your session. Unless you are working in multiple sessions and feel safe leaving your system unprotected for the time.

```
lidsadm -S -- -LIDS
```

Now with LIDS disabled you can proceed with your work.

9.3.1   An Example: Protecting a Freshly Installed Package

For this example we added a package called my_package.rpm. my_package.rpm has a configuration file in `/etc`, a binary in `/sbin`, a log is kept `/var/log/my_package.log` and stores user data in `/var/lib/my_pack age/`. `my_package.rpm` also requires *setuid* and *setgid* access. Without re-configuring LIDS this application won't function properly. Here is what needs to be done to add this package to your LIDS configuration. Issuing the following command will give you a list of the files an RPM uses. Though it won't tell you if it needs, read, write and/or append access to them.

```
rpm -qpl package_name.rpm
```

The first thing we want to do now is protect the configuration file. The configuration file never needs to be changed by the program so we can give it READ access only. If you want to make changes in the future simply disable LIDS, make your changes and enable LIDS. Here is how to protect our config file for READ only access:

```
lidsadm -A -o /etc/my_package.conf -j READ
```

Now the file is in the LIDS configuration file and set as read only. We used the *-A* option to ADD a new object. The *-o* object is the file my_package.conf and it's *-j* attribute is READ. Valid attributes are READ, WRITE, APPEND, DENY, and IGNORE.

NOTE: These are case sensitive and therefore must be written in all upper case letters.

We have successfully protected the configuration file. Next we will tackle the log file. The log file is simply a file that maintains a list of program events. The file never changes previous information and therefore can be set to APPEND only. So we issue a similar command as the one used for the configuration file:

```
lidsadm -A -o /var/log/my_package.log \
           -j APPEND
```

This command is almost the same as above except we set the log file to APPEND. Next we want to protect the user data. We want to be able to read and write to the user data, but we don't want root to have the ability to view the data, since it could be private information. This is also a secure method of protecting sensitive data from an intruder, if they gain root access. First we have to deny everybody access from the user data. There could be a slight problem if the user data directory contains dozens, maybe hundreds of files. This could be quite cumbersome typing in each file name into lidsadm. Well the lidsadm program allows you to protect a directory and everything under it. So now lets protect the directory:

```
lidsadm -A -o /var/lib/my_package/ -j DENY
```

Now everyone is denied access to that directory and everything in it. In fact, if you get a directory listing of /var/lib the my_package/ directory will not even be visible. So now it's safe. Too safe now actually. You have to give your my_package binary access to the data for it to run properly. To give the binary, and only the binary, access to the data, we can issue this command:

```
lidsadm -A -s /sbin/my_package_binary \
        -o /var/lib/my_package -j IGNORE
```

Once that is issued it gives /sbin/my_package_binary full access to everything in the /var/lib/my_package directory. In the example above we

-A added a new *-o* object but this time linked it to a *-s* subject. So now the user data is completely protected and is not hindering the usage of the my_package application.

Finally we need to protect the binary from being deleted. So we can simply set it as read only. We can use the same command that we used for the config file:

```
lidsadm -A -o /sbin/my_package_binary -j READ
```

When initially securing the system the entire /sbin directory was protected. To add /sbin/my_package_binary separately you can do what was done above or you can update all the items in the LIDS config. Doing this will add the /sbin/my_package_binary to the config

```
lidsadm -U
```

We are now left with one last problem. The my_package_binary needs *setuid* and *setgid* permissions to run properly. By default the setuid and setgid capabilities are disabled by LIDS (more concerning capabilities will be explained in the following sections). Using lidsadm you can assign capabilities to a specific file. The lidsadm command is similar to adding a file:

```
lidsadm -A -s /sbin/my_package_binary -t \
        -o CAP_SETUID -j INHERIT
lidsadm -A -s /sbin/my_package_binary -t \
        -o CAP_SETGID -j INHERIT
```

Now the /sbin/my_package_binary will inherit the setuid and setgid capabilities in the kernel giving it permission to use. The -t flag is used to tell lidsadm the object is special, or not a file in this case.

To make certain everything in your LIDS configuration is set properly issuing a:

```
lidsadm -L
```

will present you with a list of all the items in the configuration and their attributes. You must have lidsadm turned off to use this option. Now the entire package is done. Reload the config into LIDS and finally enable LIDS again:

```
lidsadm -S -- +RELOAD_CONF
lidsadm -S -- +LIDS
```

Now you are ready to go.

When LIDS is initially configured for EnGarde a script was created that contains all file attributes. This script can be run at any time to reset you back to the system defaults. Additionally you can create your own script file for any additions you make. This makes it much easier if you make a mistake and have to start over from scratch. A simple command to launch your script will put you back where you were instead of typing everything back in. If you are using the GD WebTool this is already done for you. The script can be something basic, here is a sample script using the example above:

```
#!/bin/bash
#
### LIDS configuration - 9/13/00
#
#### Configuration for my_package.rpm
#
 lidsadm -A -o /etc/my_package.conf -j READ
 lidsadm -A -o /var/log/my_package.log -j APPEND
 lidsadm -A -o /var/lib/my_package/ -j DENY
 lidsadm -A -s /sbin/my_package_binary \
           -o /var/lib/my_package -j IGNORE
 lidsadm -A -o /sbin/my_package_binary -j READ
 lidsadm -A -s /sbin/my_package_binary -o CAP_SETUID \
           -j INHERIT
 lidsadm -A -s /sbin/my_package_binary -o CAP_SETGID \
           -j INHERIT
#
#### End my_package.rpm configuration
```

You can even add this to your */etc/rc3.d/* (*/etc/rc.d/rc3.d/ for RedHat systems*)so the LIDS configuration is freshened on every boot up. Just make sure it's done before the kernel is sealed (*lidsadm -I*). More information about sealing the kernel is explained in later sections.

If this package is ever removed you will have to delete the entries. Using the script method above, delete out all the entries then *lidsadm -Z* and run all the scripts again. Otherwise you can issue a *lidsadm -D* for each file entry you have. For files with multiple entries, you only need enter it in once. Lidsadm will delete all entries for that file.

9.4   Kernel Capabilities

When a process is created it is given a set of capabilities from the kernel. These capabilities tell the process what it can and can not do. LIDS gives you the ability to alter these capabilities in the kernel. You can set the capabilities to apply to all processes or only specific processes. We saw how to apply capabilities to only specific processes previously in the *Adding an Entry* section and in the above example.

The default capabilities set that LIDS used is defined in the `/etc/lids/lids.cap` file. This file contains a list of the capabilities by name, with a number and a + or - symbol before it. A + enables the listed capability following it and a - disables it. Before each capability is a description of what the capability does. We suggest you keep the default capabilities. You can also find a list of all the capabilities and definitions at the end of this section and by just typing `lidsadm` or `lidsadm -h`. Issuing:

```
lidsadm -I
```

sets all the capabilities listed in the `/etc/lids/lids.cap` file. By default, in EnGarde Linux, the command is entered into the `/etc/rc.local` file so the kernel is sealed during boot up. When LIDS is disabled the capabilities return to their original settings and when you enable the kernel again they return to their previous state.

Earlier we set capabilities to a binary. We were actually linking a capability a process the binary creates:

```
lidsadm -A -s /path/to/binary -t -o CAP_NAME
```

All processes, however are protected from being killed by anyone but the owner of the process. This too can be avoided with the above process.

9.4.1   Capability Names and Descriptions

Here is a list of all the capabilities supported by LIDS and what their function is.

CAP_CHOWN In a system with the `_POSIX_CHOWN_RESTRICTED` option defined, this overrides the restriction of changing file ownership and group ownership.

CAP_DAC_OVERRIDE Override all DAC access, including ACL execute access if _POSIX_A
CL is defined. Excluding DAC access covered by CAP_LINUX_IMMUTABLE.

CAP_DAC_READ_SEARCH Overrides all DAC restrictions regarding read and search on files and directories, including ACL restrictions if _POSIX_ACL is defined. Excluding DAC access covered by CAP_LINUX_IMMUTABLE.

CAP_FOWNER Overrides all restrictions concerning allowed operations on files, where the file owner ID must be equal to the user ID, except where CAP_FSETID is applicable. It doesn't override MAC and DAC restrictions.

CAP_FSETID Overrides the following restrictions that the effective user ID shall match the file owner ID when setting the S_ISUID and S_ISGID bits on that file; that the effective group ID (or one of the supplementary group IDs) shall match the file owner ID when setting the S_ISGID bit on that file; that the S_ISUID and S_ISGID bits are cleared on successful return from chown(2) (not implemented).

CAP_KILL Overrides the restriction that the real or effective user ID of a process sending a signal must match the real or effective user ID of the process receiving the signal.

CAP_SETGID

- Allows setgid(2) manipulation

- Allows setgroups(2)

- Allows forged gids on socket credentials passing.

CAP_SETUID

- Allows set*uid(2) manipulation (including fsuid).

- Allows forged pids on socket credentials passing.

CATP_SETPCAP Transfer any capability in your permitted set to any pid, remove any capability in

your permitted set from any pid.

CAP_LINUX_IMMUTABLE Allow modification of S_IMMUTABLE and S_APPEND file attributes.

CAP_NET_BIND_SERVICE Allows binding to TCP/UDP sockets below 1024.

CAP_NET_BROADCAST Allow read/write of device-specific registers

CAP_NET_ADMIN

- Allow broadcasting, listen to multicast.

- Allow interface configuration

- Allow administration of IP firewall, masquerading and accounting

- Allow setting debug option on sockets

- Allow modification of routing tables

- Allow setting arbitrary process / process group ownership on sockets

- Allow binding to any address for transparent proxying

- Allow setting TOS (type of service)

- Allow setting promiscuous mode

- Allow clearing driver statistics

- Allow multicasting

CAP_NET_RAW

- Allow use of RAW sockets

- Allow use of PACKET sockets

CAP_IPC_LOCK

- Allow locking of shared memory segments

- Allow mlock and mlockall (which doesn't really have anything to do with IPC).

CAP_IPC_OWNER Override IPC ownership checks.

CAP_SYS_MODULE Insert and remove kernel modules.

CAP_SYS_RAWIO

- Allow `ioperm/iopl` and `/dev/port` access

- Allow `/dev/mem` and `/dev/kmem` access

- Allow raw block devices (`/dev/[sh]d??`) access

CAP_SYS_CHROOT Allow use of `chroot()`

CAP_SYS_PTRACE Allow `ptrace()` of any process

CAP_SYS_PACCT Allow configuration of process accounting

CAP_SYS_ADMIN

- Allow configuration of the secure attention key

- Allow administration of the random device

- Allow device administration (`mknod`)

- Allow examination and configuration of disk quotas

- Allow configuring the kernel's syslog (printk behavior domain name)

- Allow setting the domain name

- Allow setting the host name

- Allow calling `bdflush()`

- Allow `mount()` and `umount()`, setting up new smb connection

- Allow some autofs root ioctls

- Allow nfsservctl Allow `VM86_REQUEST_IRQ`

- Allow to read/write pci config on alpha

- Allow irix_prctl on mips (`setstacksize`)

- Allow flushing all cache on m68k (`sys_cacheflush`)

- Allow removing semaphores

- Used instead of *CAP_CHOWN* to chown IPC message queues, semaphores and share memory

- Allow locking/unlocking of shared memory segment

- Allow turning swap on/off Allow forged pids on socket credentials passing

- Allow setting read-ahead and flushing buffers on block devices

- Allow setting geometry in floppy driver

- Allow turning DMA on/off in xd driver

- Allow administration of md devices (mostly the above, but some extra ioctls)

- Allow tuning the ide driver Allow access to the nvram device

- Allow administration of apm_bios, serial and bttv (TV) device

- Allow manufacturer commands in isdn CAPI support driver

- Allow reading non-standardized portions of pci configuration space

- Allow DDI debug ioctl on sbpcd driver

- Allow setting up serial ports

- Allow sending raw qic-117 commands

- Allow enabling/disabling tagged queuing on SCSI controllers and sending arbitrary SCSI commands

- Allow setting encryption key on loopback file system

CAP_SYS_BOOT Allow use of `reboot()`

CAP_SYS_NICE

- Allow raising priority and setting priority on other (different UID) processes

- Allow use of FIFO and round-robin (realtime) scheduling on own processes and setting the scheduling algorithm used by another process.

CAP_SYS_RESOURCE

- Override resource limits. Set resource limits.

- Override quota limits.

- Override reserved space on ext2/ext3 file system

- NOTE: ext2/ext3 honors fsuid when checking for resource overrides, so you can override using fsuid too

- Override size restrictions on IPC message queues

- Allow more than 64hz interrupts from the real-time clock

- Override max number of consoles on console allocation

- Override max number of keymaps

CAP_SYS_TIME

- Allow manipulation of system clock

- Allow irix_stime on mips

- Allow setting the real-time clock

CAP_SYS_TTY_CONFIG

- Allow configuration of tty devices

- Allow `vhangup()` of tty

A QUICK START GUIDE

This appendix is intended to give an overview of the functions of the Guardian Digital WebTool. After reading this appendix, the reader should be able to perform the steps required to set up a domain to receive mail, configure DNS services, and serve Web pages. If your EnGarde system will not be used to perform all of the functions listed above, it is especially important that you read the User Guide and have a full understanding of each of the services you will be configuring.

Before following the example below, EnGarde should have already undergone initial configuration and be plugged in and operating on a network. Information regarding the initial configuration can be found in *Section 3 Installing EnGarde* on page 12.

To obtain a fast and most accurate setup, follow the steps in the described order. Once you have successfully completed each step, proceed in order to the next step. There are four primary steps required to configure EnGarde:

1. Configure the network interface

2. Configure the DNS Server

3. Configure the Mail Server

4. Configure the Web Server to prepare for normal and secure websites

After the initial configuration of your EnGarde Secure Professional system, the basic system and networking functions are operating correctly and is ready to configure a sample store. We will be configuring our example EnGarde system to use the following initial values entered when EnGarde was configured:

Hostname: myserver

Domain Name: mydomain.com

IP Address: 192.168.1.70

Netmask: 255.255.255.0

Gateway: 192.168.1.1

Primary DNS Address: `192.168.1.70`

Secondary DNS Address: `192.168.1.60`

In this example, we will be creating the domain `engardelinux.com` that will be hosting our DNS, routing mail, and serving web pages.

A.1   Network Interfaces

Before any interfaces are created you will need to know the following:

- Each SSL-based website requires its own IP address. If more SSL-based websites are to be served, then a new interface must be created on another IP address for each website.

- There can be many normal websites on the same IP address, given a *Name Virtual Host* defined in the Web server. See the *Section* 4.3 *Virtual Host Management* on page 56 in the *User Guide* for more information on *Name Virtual Hosts*.

Example:

In the WebTool, click on *System Management*, and then click on *Network Configuration*. There will already be an interface defined as:

We want to set up a separate IP address for `www.engardelinux.com`, since we will be creating a *Secure Web Server* on it. Click on *Add a New Interface* to do this. We are now prompted for our information, at which point we enter:

IP Address: `192.168.1.71`

Netmask: `255.255.255.0`

After clicking the *Create* button the *Persistent Interfaces* screen will look like:

	IP Address	Hostname
[Edit]	192.168.1.70	myserver.mydomain.com
[Edit]	192.168.1.71	< Not Yet Defined >

Add a New Interface

We have now successfully configured our network interface.

A.2   DNS Server

The DNS Server is the mechanism that provides name to IP address, and IP address to name mappings. It also provides the information necessary for mail to be properly routed. DNS was created because IP addresses are often hard to remember. DNS is used to map that address to a name, which is much easier to remember.

When typing `http://www.guardiandigital.com` into a Web browser, for example, the DNS server translates the host name (`www.guardiandigital.com`) into the IP address associated with `www.guardiandigital.com`. The browser then sends the request to that IP address and responds with the information available at that address.

DNS contains a number of unique characteristics about each host. Each characteristic forms a 'record' in the database that stores the DNS information. DNS "zones" are regions of IP addresses or names for which a particular organization is responsible.

Address Records This is a record that provides a host name to be assigned to an IP address. All host names are associated with an IP address.

Name Server Records This is a record that defines what name servers are responsible for the zone. In most cases, this will be the same as the hostname of the machine. Do not alter these records unless you have an explicit reason to.

Name Alias Records This is a record which provides an "alias" for a pre-existing host name. There may be multiple aliases for a single host name.

Mail Server Records This is a record which provides the information necessary to correctly route mail to correctly deliver electronic mail. Multiple e-mail

servers may be defined for the same domain, each with a differing priority. Servers defined with a lower number have a higher priority and mail will be delivered to these hosts first.

Example:

Because we are creating a new domain (`engardelinux.com`), we must create a new forward zone for it. Before EnGarde can be configured to provide DNS for this domain, it must have been listed among the list of authoritative name servers for this domain.

>From the *System Management* menu, select *DNS Management*. The next step will be to create a new master zone. Click on the *Create a New Master Zone* link.

Leave the *Forward (Names to Addresses)* button checked since that is the type of zone to be created. Keep the default value of *Master server*. The rest the input looks like:

Domain name: `engardelinux.com`

Email Address: `administrator@engardelinux.com`

Leave the *Allow transfers from...* set to *Allow None*, and the *Allow queries from...* set to *Allow Any*. For more information on these fields please refer to the full manual.

Click on the *Create* button to see the new zone in the zone listing. To add the records for our example, click on the *engardelinux.com* link.

Address Records

Hostname: `www.engardelinux.com`
Address: `192.168.1.71`

Hostname: `mail.engardelinux.com`
Address: `192.168.1.71`

Name Alias Records

Alias: `sales.engardelinux.com`

Real Name: `www.engardelinux.com`

Mail Server Records

Mail Server: `mail.engardelinux.com`

Priority: `10`

At this point we have successfully created `www.engardelinux.com` and `mail` `.engardelinux.com` to go to `192.168.1.71`.

We have now successfully configured the DNS records for our sample domain.

A.3   Mail Server

The mail server provides the mechanism to deliver e-mail to a recipient on the Internet. When an e-mail is sent, the mail server is instructed to deliver the message to the remote mail server responsible for the recipient's domain.

Example:

To configure e-mail for our new domain, we must create a new Mail Domain. From the *System Management* section select *Mail Server Management*. Then select *Domain Management*.

We want to *Create [a] New Domain* with the following values:

Domain: `engardelinux.com`

Postmaster: `ryan`

This assumes that there is a user named *ryan* on the system. Now EnGarde has been configured to receive mail for `engardelinux.com`. The local user *ryan* has been defined as the Postmaster. More information on the "Postmaster" account is available in *Section* 4.4.4 *Mail Server Managemen*t on page 87.

Once the mail domain is created, individual user accounts can be added by clicking on the `engardelinux.com` link:

Example 1:

    E-Mail Username: `administrator`

    Recipient: `christi`

Example 2:

    E-Mail Username: `info`

    Recipient: `christi`

Example 3:

    E-Mail Username: `webmaster`

    Recipient: `ryan`

Example 4:

    E-Mail Username: `sales`

    Recipient: `fred@guardiandigital.com`

Here four e-mail addresses are defined. The following table shows the destination of various e-mail addresses according to the examples defined above:

Mail Sent To:	Final Recipient:
`administrator@engardelinux.com`	`christi`
`info@engardelinux.com`	`christi`
`webmaster@engardelinux.com`	`ryan`
`sales@engardelinux.com`	`fred@guardiandigital.com`
`ryan.maple@engardelinux.com`	`ryan`

We have now successfully configured our Mail Server.

A.4   Web Server

The Web Server is the mechanism for serving websites. There are two types of websites: *normal* and *secure*. Secure websites utilize SSL encryption to provide

security for sensitive applications such as e-commerce. Normal websites are simply sites that do not utilize SSL.

Secure websites require two things: a certificate and a key. It can be thought of in the following context:

- the certificate is what verifies your identity (authentication)

- the key is what provides the security (encryption)

The certificate and key are also tightly tied into each other; they are a matching pair.

The first time a user connects to a secure site, their browser will store the certificate. Every subsequent time the user connects to the site it verifies that the certificate is the same to ensure a secure connection. This provides the encryption portion of the process.

For more information on certificates please refer to the full User Guide.

Example:

To configure the Web server for our new domain, we must set them up in *Section 4.3 Virtual Host Management* on page 56.

To create the normal site, go to *Virtual Host Management*, and select *Create a Virtual Host*. We use the following values:

Address: `192.168.1.71`

Administrator E-Mail: `webmaster@engardelinux.com`

Server Name: `www.engardelinux.com`

Webmaster: `ryan`

For *Group*, we want to first *Create [a] Group* named *engardeweb*, and then select it.

Group: `engardeweb`

If a database is necessary for this site, then we check the *Create a database for this site* box and enter in the values:

Username: `engardeweb`

Password: `e!nGa#rDe`

We have now successfully created the normal website.

Likewise, to create the secure site, go to *Virtual Host Management*, and select *Create an SSL Virtual Host*. We use the following values:

Address: `192.168.1.71`

Administrator E-Mail: `webmaster@engardelinux.com`

Server Name: `www.engardelinux.com`

Webmaster: `ryan`

Group: `engardeweb`

We have now successfully created the secure website.

Once this is done, the following directories for the normal site will be created:

```
/home/httpd/www.engardelinux.com.com-80/cgi-
bin
/home/httpd/www.engardelinux.com-80/html
/home/httpd/www.engardelinux.com-80/logs
```

And the following directories for the secure site:

```
/home/httpd/www.engardelinux.com-443/cgi-bin
/home/httpd/www.engardelinux.com-443/html
/home/httpd/www.engardelinux.com-443/logs
/home/httpd/www.engardelinux.com-443/ssl
```

Once the above steps have been completed, EnGarde is ready to serve webpages for the following sites:

```
http://www.engardelinux.com/
https://www.engardelinux.com/
```

The next step is to populate your sites with content. For more information on this and the many other aspects of the WebTool, please refer to the User Guide.

B ADVANCED INSTALLER ISSUES

B.1 Boot disk creation

If your PC does not support the ability to boot from a CD-ROM then you must create a boot floppy. A boot floppy simply contains the same boot image that is on the CD.

To create a boot floppy have a blank floppy available and the EnGarde Secure Professional CD-ROM in the drive, and if in a Linux system, mounted as well.

B.1.1 Creation on a Linux based system

The boot image is located on the CD in `~/boot/boot.img`. Type the following command in a shell to create a boot disk:

```
# dd if=/mnt/cdrom/boot/boot.img of=/dev/fd0 bs=1k
```

The above command assumes the CD is mounted in `/mnt/cdrom`, change this if necessary. Once you have been returned to the prompt the disk is ready for use.

B.1.2 Creation on a DOS based system

Included on the CD-ROM are DOS utilities for creating a boot disk. Inside of `x:\dosutils` you will find a program called `rawrite.exe`. This will write the image to the floppy disk.

NOTE: Replace `x:` throughout this example with the assigned drive letter of your CD-ROM drive.

From a prompt type the following:

```
C:\> x:\dosutils\rawrite.exe -f x:\boot\boot.img -d a:
```

Once this has completed your boot floppy is ready for use.

B.2 Rescue mode

EnGarde Secure Professional includes a rescue mode in the installer. Rescue mode will boot up a working Linux system off of the EnGarde CD-ROM and allow you to trouble shoot your system.

Rescue Mode can be accessed by typing in *rescue* or *linux rescue* at the LILO boot prompt. Rescue mode requires that the EnGarde CD-ROM be in the drive regardless if you are booting from the CD or a boot floppy. The rescue system is located on the CD.

WARNING: Rescue mode is for experienced Linux users only. An existing En-Garde installation can possibly be damaged if used improperly.

Once the system boots you will have a working Linux system which includes many programs to help you recover your system.

To reboot from rescue mode simply make certain all your hard drives have been unmounted and simply press CTRL-ALT-DEL and remove all bootable media from the machine.

B.3   Automatic partition scheme

When selecting *Automatic Partitioning* the installer will partition up your drive with predefined rules. Here is how the installer decides how to break your drive up:

- '/' (root) will be 25% of the drive but no less than 320MB and no greather than 2048MB

- The swap partition is 5% of the system drive but will not be less than 32MB and no greater than 256MB.

- /var and /home will them be 50% each of the total remaining space.

For example, if we have a 20Gb drive (20012MB) the partitions will look like this:

```
/        2048MB

swap     256MB

/home    8854MB

/var     8854MB
```

These numbers are determined as follows:

```
/        20012 * .25 = 5003MB. 5003 > 2048MB.

swap     20012 * .05 = 1000.6MB. 1000.6 > 256MB.

/home    20012 - (2048 + 256) = 17708 * .50 = 8854

/var     same as /home
```

C GENERAL LINUX

C.1 Introduction

In this section we will discuss some basic Linux knowledge for administering En-Garde from the console or an SSH connection. This section is more for advanced users. You have to be careful, you can corrupt the system configuration resulting in improper operation of your EnGarde system.

C.1.1 Root Access on Your EnGarde System

su is a small program that gives you the ability to login as the root user from a remote connection. To help increase security you are prevented from running *su*. The only ways to gain root access is to either login as root from the console or make an SSH connection to EnGarde as the root user.

All logins via *SSH*, both root logins and normal user logins are logged in `/var/log/syslog` and are filtered into `/var/log/audit/ssh_authorization.log`, `/var/log/audit/su_logins.log`, and `/var/log/audit/su_failed.log`. You can find console logins in the `/var/log/audit/pam.log` which will contain all successful and failed login attempts from the console.

C.2   Basic Bash Commands

Bash, or the Bourne Again Shell, is the successor to *sh*. Bash is the default system
shell you will be using to interface with EnGarde when you login via SSH or the
console. Here we will cover some basic commands for moving around the system
and doing some minor work. If you will be doing most of your editing from the
command line we highly recommend picking up a book on using bash or general
Unix commands.

NOTE: You will find `/bin/sh` on your system. It is really a link to `/bin/bash`.
            This is done for compatibility reasons.

C.2.1   Moving Around the System

When you first login you will be sitting in your home directory. Most likely
`/home/username/`. You can get a listing of the directory contents by typing:

```
$ ls
```

or for a long view of the listing with time stamps, file permissions and file owner-
ships type:

```
$ ls -l
```

You can move from directories by typing

```
$ cd directory-name
```

cd by itself will bring you back to your home directory.

Directories are referenced with a slash (/). / being the root directory. So to go
to the */etc* directory you simply type

```
$ cd /etc
```

to reference the current directory we use a single period, '.' and to reference the
previous directory we use two periods, '..'. So if you are in your home directory
and you want to go to a different users directory you can type:

```
$ cd images/different-user
```

which is equivalent to:

```
$ cd /home/different-user
```

At any point using the TAB key after typing a few characters in at the bash prompt will make bash fill in the rest of the file. or directory name that matches what you have typed. If there is more than one match, tap the tab key twice and it will list all the matches.

C.2.2   File Manipulation

There are many ways to alter files on your system. You can copy, delete, move, change attributes etc. Here is the three basic file manipulation commands, cp, rm, and mv -> Copy, remove and move. They are used as follows:

```
$ cp file1 file2
ex: $ cp /home/nick/new_httpd.conf /etc/httpd/conf/httpd.conf
$ rm file
ex: $ rm /home/nick/new_httpd.conf
$ mv file1 file2
ex: $ mv /home/nick/new_httpd.conf /etc/httpd/conf/httpd.conf
```

You also have control over the attributes and ownership of a file. Running *chown* and *chgrp* you can change the files ownerships:

```
$ chown nick *.html
$ chgrp nick *.html
```

The above two commands will give user nick complete ownership over every html file in the current directory. You can shorten the above command by typing:

```
$ chown nick:nick *.html
```

This changes both the ownership and group in one shot. You can change the file permissions using the *chmod* program. By typing:

```
$ chmod 644 *.html
```

That will change the access to read/write by the owner and read only by users in the specified group and all users. There are many more options, too many to list here, *chmod* can use.

C.2.3   Editing a File

You basically have two options for file editing from the console, Vi and Pico.

Vi has the most difficult learning curve but is the most powerful editor. Pico is much easier to learn. All the commands are laid out in front of you. Pico, however can have some strange effects on files and is not nearly as powerful as the other two editors.

EnGarde comes with Vi and Pico installed on it. To load the Vi editor simply type:

```
$ vi fileToEdit
```

To start the Pico editor type:

```
$ pico fileToEdit
```

If you don't enter a filename it will start by editing a blank document.

We recommend using Vi if you will be doing most of your editing from the console. If you don't have experience with *vi* you'll want to use one of the many resources as it's use may not be immediately obvious.

C.3   File System Structure

The EnGarde Linux system is designed with the file system standards in mind. Here is a brief breakdown of the directories and there descriptions (taken from Filesystem Hierarchy Standard - ver2.1):

```
/ - the root directory
|-bin     Essential command binaries
|-boot    Static files of the boot loader
|-dev     Device files
|-etc     Host-specific system configuration
|-home    User home directories
|-lib     Essential shared libraries and
|         kernel modules
|-mnt     Mount point for mounting a
|         filesystem temporarily
|-root    Home directory for the root user
|-sbin    Essential system binaries
|-tmp     Temporary files
|-usr     Secondary hierarchy
|-var     Variable data
```

This is just a brief summary of the main root file system. For more detailed information you can download the Filesystem Hierarchy Standard from `http://www.pathname.com/fhs/`.

C.4   Services and Daemons

Linux has the ability to start and stop services and daemons on the fly. A service is generally something like POP3 or an FTP server and are managed using files in the /etc/inet.d/ directory. You can also have services ran from the init.d scripts. Here are a few commands with their results:

```
$ /etc/init.d/crond start
Starting crond:                    [ OK ]
$ /etc/init.d/d stop
Shutting down crond:               [ OK ]
$ /etc/init.d/crond restart
Shutting down crond:               [ OK ]
Starting crond:                    [ OK ]
$ /etc/init.d/crond status
crond (pid 18529 18525 18522) is running
```

Not all commands in this directory have the above options. To get a list of what each one can do, type the filename by itself.

This is primarily used if you need to shutdown a daemon for maintenance or other reasons. Remember, when you make modifications to configuration files for a daemon, you generally have to restart that daemon before the changes can take effect.

C.5   Groups and Users

File and directory permissions are the basic means for providing security on a system. They are also the last line of defense against an unauthorized user reading or modifying information that does not belong to them. A properly configured system contains files and directories which are only accessible to the users in which were authorized to access those files and directories. The set of rules that a file or directory is given to tell it who can and can't access it are known as permissions. These file and directory permissions are assigned by both user and group.

Each file and directory has three sets of permissions associated with it. It gives permissions to *owner*, *group* and *other*. Below is the result of a sample directory listing produced by executing `ls -l`, displayed with each field broken down:

```
        |----1----|-2--|---3----|----4-----|---5--|-----6------|-----7-----|
        -rw-r--r--   1 nick      users       6619   Oct 24 15:57 README
```

Field 1:   Permissions for this file. We will break down these nine file permission settings in the next section.

Field 2:   Number of hard links to this file or directory. These links can be directories.

Field 3:   Owner of the file. The users user name is displayed, if no user name is associated with the owner then the user ID number is displayed.

Field 4:   The group to which the file belongs. A group name will be displayed here, if no group name is associated with the ID then the ID number is displayed.

Field 5:   This is the size of the file in bytes.

Field 6:   The date of the last time the file was modified.

Field 7:   The name of the file.

There are three options for file permissions. Read (r), write (w) and execute (x). These three options can each be assigned to the *user*, *group* and *other* attributes of each file and directory. We can break down field one above as follows:

```
1222333444
-rw-r--r--
```

1. Special Flag

2. Owner permissions

3. Group permissions

4. Other permissions

We have S as a special attribute. Here is a list of special attributes:

- d - Directory

- s - socket

- b - block special file (IE: `/dev/hda`)

- c - character special file (IE: `/dev/tty`)

- l - sybolic link

- p - named pipe

Next we have the owner of the file, followed by the group and finally the other. Each one can have their own set of read, write and executable permissions.

D   FIREWALLS AND PROXY SERVERS

D.1   Configuring a Firewall or Proxy Server

A firewall is a system designed to keep everything behind it safe from the outside world. It scans incoming connections and determines whether or not the connection matches one of a list of pre-defined access control rule, accepts or rejecting the connection.

If your EnGarde system will be positioned behind firewall you will need to configure your firewall to allow EnGarde access to the outside world. Below are a list of ports and what they are. You may not have all of the listed ports opened on your EnGarde system if you don't have it configured to. For example, if your EnGarde system is not a DNS server you will not have the DNS port 53 opened.

22/tcp	This is the SSH port. If you want to allow anyone from outside to SSH into your machine you must open this port
25/tcp	This is the SMTP service. If this machine will be receiving e-mail this port must be available.
53/tcp&udp	This is the DNS service. You will need to have this opened. Configuring DNS to work through a firewall or proxy server can be difficult and it is recommended to refer to your firewall manual for complete instructions.
80/tcp	If EnGarde is going to be a Web server you will need to enable access to this port.
443/tcp	If EnGarde is a Web server and will be hosting a secure site you will need to open this port to support SSL
993/tcp	If EnGarde will be offering Secure IMAP you will need to have this port open.
995/tcp	Secure POP3 will be available from this port if EnGarde is running it.
1022/tcp	This is the user password changer portion of the GD WebTool. If you want to give outside users to availability to change their own password via the GD WebTool you will need to open this port up.

1023/tcp   This is the actual GD WebTool for the administrator. If you will be
             administrating this from outside you will need to open the port.

For more information about firewalls there are many books and on-line documentation. Refer to your firewall documentation for specific instructions on how to permit these services through your firewall. Additionally, here are a few references:

- Zwicky, Cooper & Chapman. Building Internet Firewalls, June 2000. Copyright O'Reilly & Associates, Inc. 2000.

- Mark Grennan, mark@grennan.com. Firewall and Proxy Server HOWTO, Feb. 26, 2000. Copyright Mark Grennan, 2000.
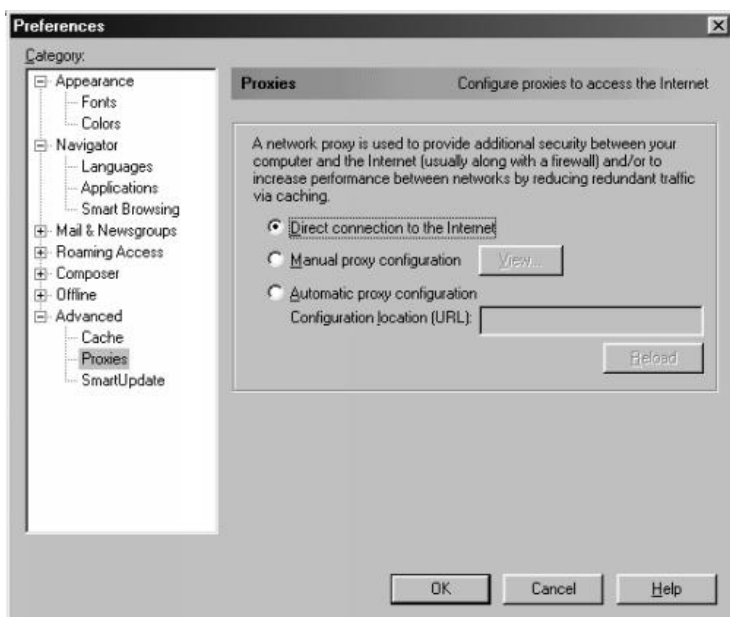
D.2 Disabling Proxy Settings in Your Browser

You will need to disable proxy and firewall settings in your browser in order to access the inital configuration tool on EnGarde. Directions are given below for both Netscape Navigator and Internet Explorer.

D.2.1 Netscape Navigator

To disable the proxy settings in Netscape Navigator you will need to be at the main Netscape Navigator window. Click the *Edit* menu button and then select *Preferences* from the pull-down menu.
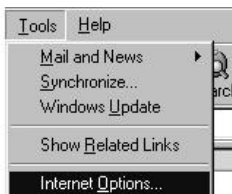

You will then be brought to the *Preferences* menu. By clicking on the *Advanced* option in the menu "tree" on the left will bring up the *Proxy Settings*.

Click the radio button labeled *Direct connection to the Internet* and then click *Ok*.
Your Netscape browser is now ready to connect to your EnGarde system.

D.2.2   Internet Explorer

To disable the proxy settings in Internet Explorer you will need to be at the main
Internet Explorer window. Click the *Tools* menu button and then select *Internet
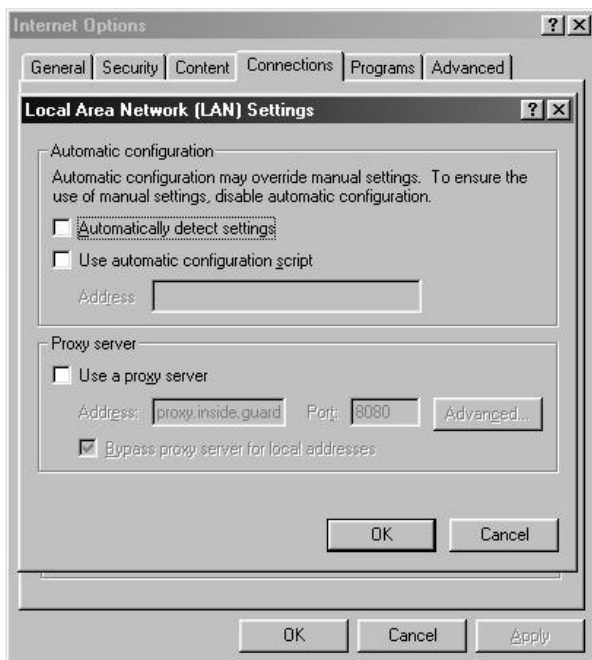Options* from the pull-down menu.

Once you select *Internet Options* you will be presented with the Internet Options

dialog box. At the top of the box there are a list of tabs, select *Connection*. From the *Connection* section click the *Lan Settings* button.

After clicking the *Setup* button the proxy information will be displayed. You want to turn off all your proxy server settings so you have to make sure all the checkboxes are NOT checked. Once this is done click the *OK* button to finish.

You are now ready to connect to your EnGarde system with Internet Explorer.

E CERTIFICATES

E.1 General Certificate Information

Here we will just briefly cover some basic certificate information you may need to know to get your certificates properly working.

A new certificate is only valid for 365 days, or 1 year. After this period you must get a new certificate. If you have a signed certificate you have the option to renew that certificate, which usually requires a fee.

E.1.1 Getting a Certificate Signed

The two most common certificate companies are Verisign and Thawte. To get a certificate signed, generate a CSR as described in *Certificate Management* found in *Section 4.3* and follow their directions to send it to the appropriate CA.

They will then request proof of your right to use the certified organization name (Articles of incorporation), proof of your registration of the domain name you will be using (from the InterNIC whois database), to obtain your domain name details go to:

```
http://rs.internic.net
```

And finally a letter of authorization from an agent of your company or organization.

Once everything is authorized they will send you back a signed certificate. Please read their Web sites:

```
http://www.verisign.com
http://www.thawte.com
```

for detailed information on submitting a certificate to be signed or go directly to their registration pages:

```
http://digitalid.verisign.com/server/enrollIntro.htm
http://www.thawte.com/certs/server/request.html
```

If you get a certificate signed by a smaller Certificate Authority, Netscape and Internet Explorer may bring up a warning that it does not recognize the CA. This may make some users uncomfortable and insecure about using your site. However, one of these CAs can provide you with a signed certificate at a much reduced cost.

E.1.2   Certificates, IP and Virtual Host Issues

A certificate is bound to a domain name regardless of the IP address. Therefore if you register a certificate you will register it under your domain name. Unfortunately due to current protocol restrictions you can only have one certificate per IP address.

Using a separate IP for each domain name located on your EnGarde system will give you the ability to assign a separate certificate to each domain.

E.2   Accepting an Unsigned Certificate

During the initial login during the configuration of your EnGarde system and/or when connecting to the GD WebTool you will be prompted with the following screen:

Your browser will ask you if you want to accept the certificate attached to your EnGarde system. The reason for this is Guardian Digital has signed the certificate and is not a Certificate Authority (CA) such as Verisign and Thawte. Having this certificate signed by a CA is not necessary since you can verify that you are connecting to your own EnGarde system.

You will want to accept this certificate. Click the *Next* button to continue.

This next screen will display brief information concerning the certificate. There is a button you can click, *More Infor...* for detailed information concerning the certificate. Click *Next* to continue.

Now you will be asked in what way you want to accept this certificate. You have three options here. The first option will only accept the certificate for the current session. So when you shut your browser down you will be prompted with the same screens the next time you try to login to the GD WebTool.

The second option will tell your browser to never accept the certificate. This will lock you out of GD WebTool.

Finally the third option will accept the certificate until it expires. When it expires and a new certificate is put in it's place you will be prompted again with these same menus.

If you will be doing your administration via the GD WebTool on the current machine it is recommended you select *Accept this certificate forever (until it expires)* option. Once you have made your decision select the *Next* button.

This fourth screen will inform you of the possibility of fraud and insecurity when using an unsigned certificate. Since you know EnGarde Linux and the certificate both came from Guardian Digital you can be certain your connection and data will be secure.

This is the final step and will inform you of your decision to accept the certificate and verify your options. Click *Finish* to fully accept the certificate and enter the GD WebTool.

F GLOSSARY

attributes (ext2fs-specific) In addition to standard Unix permissions, the ext2 file system contains additional attributes, which the file system driver honors whenever the file is accessed or modified. Attributes are set or unset by the CHATTR command, and it is common to refer to the bits set by the name. The "immutable" bit is particularly popular among system administrators trying to protect critical files from unintentional destruction by an inattentive ROOT user.

authentication The process of knowing that the data received is the same as the data that was sent, and that the claimed sender is in fact the actual sender.

backup (or archive) Both of these terms are used as nouns and verbs. The noun form refers to any copy of a set of files (and the *meta-data* associated with them) on some form of removable media. The verb form refers to any process of creating such a set. An extra copy of a set of files to non-removable storage is sometimes referred to as "*a backup*"– but this is more precisely referred to as "*replication*" or "*mirroring*" or (in some cases) "*version control*"

bastion host A computer system that must be highly secured because it is vulnerable to attack, usually because it is exposed to the Internet and is a main point of contact for users of internal networks. It gets its name from the highly fortified projects on the outer walls of medieval castles. Bastions overlook critical areas of defense, usually having strong walls, room for extra troops, and the occasional useful tub of boiling hot oil for discouraging attackers.

broadcast The broadcast address is a special address that every host on the network listens to in addition to its own unique address. This address is the one that datagrams are sent to if every host on the network is meant to receive it. Certain types of data like routing information and warning messages are transmitted to the broadcast address so that every host on the network can receive it simultaneously. There are two commonly used standards for what the broadcast address should be. The most widely accepted one is to use the highest possible address on the network as the broadcast address. An

example on an internal network would be 192.168.1.255. Every host on the network must be configured with the same broadcast address.

buffer overflow Common coding style is to never allocate large enough buffers, and to not check for overflows. When such buffers overflow, the executing program (daemon or set-uid program) can be tricked in doing some other things. Generally this works by overwriting a function's return address on the stack to point to another location.

denial of service An attack that consumes the resources on your computer for things it was not intended to be doing, thus preventing normal use of your network resources for legitimate purposes.

DHCP See *Dynamic Host Configuration Protocol.*

DNS See *Domain Name Server.*

Domain Name Server The Domain Name System (DNS) is the software that is responsible for converting hostnames into numbers that computers can understand. For example, the name www.guardiandigital.com corresponds to the host IP address 63.87.101.80 and vice versa. The DNS server, sometimes called a name server, is the process that runs on EnGarde awaiting incoming name service requests.

dual-homed host A general-purpose computer system that has at least two network interfaces.

Dynamic Host Configuration Protocol Also known as DHCP, is a protocol for assigning dynamic IP addresses to devices on a network. DHCP simplifies network administrative work because the software keeps tracks of IP addresses as opposed to the administrator.

EXT2 Is the main filesystem the Linux operating system uses on its storage devices.

EXT3 A filesystem based on the EXT2 filesystem that includes journaling capabilites.

filesystem The filesystem manages files contained on a storage device so that the operating system may interact with them. The most common filesystem in Linux is Ext2.

firewall A component or set of components that restricts access between a pro-
tected network and the Internet, or between other sets of networks.

forward zone A forward zone contains a listing of the hostnames in that zone
with their corresponding IP addresses. A reverse zone represents address-
to-domain mapping, such as `63.87.101.80` to `www.guardiandigital`
`.com`.

forwarder A forwarder is used for name servers that may not necessarily be
directly-connected to the Internet. This may be due to being behind a fire-
wall, or inside of a corporate network. Forwarders will instead only query
a specified additional name server for its DNS information.

FQDN See *Fully-Qualified Domain Name*.

Fully-Qualified Domain Name Domain names reflect the domain hierarchy. Do-
main names are written from most specific (a host name) to least specific
(a top-level domain), with each part of the domain separated by a dot '.'.
A fully qualified domain name (FQDN) starts with a specific host and ends
with a top-level domain. An example of this could be:

Name	Type
`engarde.guardiandigital.com`	FQDN
engarde	Machine Name
`guardiandigital.com`	Domain Name
`com`	Domain

full backup This is probably the most confusing term that relates to the subject
of backups. It often does not mean "*comprehensive*." A "*full*" backup does
not necessarily mean that it includes every file on a whole system. "*Full*"
in those cases means "*including all files in a given data set without regard
to previous backups*." In other words, it means "*not incremental*" and not
"*differential*." It is better to use the phrase "*level zero*" to make this distinc-
tion.

GNU GNU's Not Unix, a recursive acronym. This is the name of a project started
by Richard M. Stallman, and is the mission of the FSF (Free Software Foun-
dation), which he founded. The purpose of the GNU project is to produce

a "free" operating system and suite of applications, utilities, and programming tools that are non-proprietary and unencumbered.

GPL To protect the GNU project software from being appropriated for proprietary use by hardware vendors, the Free Software Foundation released their software under the GPL or General Public License.

hard link An entry in a directory that contains a pointer directly the the inode bearing the file's *meta-data*. All non-symlink directory entries are " *hard links*."

host A computer system attached to a network.

host key A key the host will store locally and used for authentication when a user key, stored on the users system, is passed to it. If both keys are valid then both the host and user. Usually associated with SSH.

IDE *See Intelligent Drive Electronics*.

Intelligent Drive Electronics An interface for mass storage devices that have the controller integrated into the disk. Also refered to as IDE for short.

Internet Message Access Protocol A protocol for retrieving e-mail from a server. Similar to POP3 but instead of downloading messages to the local machine IMAP's default is to work on the server.

IP spoofing IP Spoofing is a complex technical attack that is made up of several components. It is a security exploit that works by tricking computers in a trust relationship into thinking that you are someone that you really aren't. There is an extensive paper written by daemon9, route, and infinity in the Volume Seven, Issue Forty-Eight issue of Phrack Magazine.

ISO9660 The most common file system found on CD-ROMs.

Kernel Unix systems have a kernel that provides a system call interface (including IOCTL() I/O device control interface) to allow programs to interface directly with hardware and files. The Linux kernel provides file systems, networking support for TCP/IP and other protocols, and device drivers. These can be built into a kernel "*statically*" or as loadable modules.

LIDS See *Linux Intrusion Detection System*.

Linux Intrusion Detection System The Linux Intrusion Detection System allows fine tuning of control over resources and file permissions. For detailed information concerning LIDS and using LIDS please read Section 9.

loadable modules Portions of kernel code that have been compiled separately and that can be loaded during normal operation using *modprobe* or *insmod*. If you have LIDS running it seals the ability to load modules after the system has booted. You must shut LIDS off first, then load your module(s). Information on controlling LIDS can be found in Section 9.

journaling Journaling is a method used to preserve data when it is written to a storage device. This greatly increases recovery time in the event of a system crash.

mount A storage device containing a device can not be accessed by a Linux system until it is mounted. The process of mounting allows the system to make a common "reference" to this filesystem. This is done by mounting a filesystem to an empty directory. The filesystem will then be contained within that directory.

non-repudiation The property of a receiver being able to prove that the sender of some data did in fact send the data even though the sender might later deny ever having sent it.

Open Source Programs for which the original source code is available, for which relatively permissive opportunities to modify the code and share the results with others exist, and which are developed by people whose primary means of communication with each other is the Internet.

OpenSSH An Open Source version of Secure Shell.

ownership The user (UID) and/or group (GID) that is associated with a file, directory, process, or process group.

packet The fundamental unit of communication on the Internet.

packet filtering The action a device takes to selectively control the flow of data to and from a network. Packet filters allow or block packets, usually while routing them from one network to another (most often from the Internet to an internal network, and vice-versa). To accomplish packet filtering, you set up rules that specify what types of packets (those to or from a particular IP address or port) are to be allowed and what types are to be blocked.

partition Before a storage device such as a hard drive can be used by the system it must be partitioned. A partition is a portion of the whole drive. It defines the boundries in which the filesystem can manage. A filesystem can not be placed on a storage device without a designated partition.

partitioning See *partition*.

perimeter network A network added between a protected network and an external network, in order to provide an additional layer of security. A perimeter network is sometimes called a DMZ.

pid Process identifier. A number used by the kernel to keep track of the system-level resources necessary to switch between this process and others running on the system. It is easily visible to a system administrator by use of the *ps* command. In the GD WebTool, Section 4, you will find detailed instructions on viewing and deleting processes via the WebTool.

pptp See *Point-to-Point Tunneling Protocol.*

protocol A predefined standard for transmitting data between two devices.

proxy server A program that deals with external servers on behalf of internal clients. Proxy clients talk to proxy servers, which relay approved client requests to real servers, and relay answers back to clients.

Point-to-Point Tunneling Protocol A secure protocol for transmitting data necessary for a Virtual Private Network (VPN) over the Internet.

Post Office Protocol A protocol for retrieving e-mail. Also refered to as POP3 (version 3), it downloads all new e-mail messages from the server and stores them locally on a users machine.

reverse zone See *forward zone*.

root Root is the "superuser" of the system. Generally the system administrator will login with root privileges to administer the system. You can not login remotely as root, only from the console. It is not recommended to login as root unless you need to since accidental errors can be easily made.

samba A client/server for non-Windows based system integration into Windows File Sharing and Printing system.

SCSI See *Small Computer System Interface.*

Secure Shell A secure shell is a telnet type connection made to a remote host. This connection is protected with SSL 3DES 128bit encryption. Secure shell is also known for short as SSH. It is pronounced S-S-H.

Secure Socket Layer Is a protocol designed by Netscape Communications that provides encrypted communications for private documents via the Internet. SSL works by use of a public/private key system for exchanging session keys.

shared libraries Shared libraries are object files that are dynamically linked to executable binary programs. Under Linux, shared libraries can be stored in a number of directories (usually listed in /etc/ld.so.conf). Shared libraries typically include files under /usr/lib. If the shared libraries are deleted or become damaged, or of the /etc/ld/so.cache file is corrupted, then programs that rely on them will fail to execute. Almost all normal programs on a system rely on glibc.

signal Under Unix and Linux, the signal is the most fundamental and common form of interprocess communications (IPC). It is also the basis for "event-driven" programming under these systems. Each Unix implementation defines a set of signals that area associated with various asynchronous events, such as a terminal sending an "interrupt request" (SIGINT) or a change in window size (SIGWINCH).

SIMAP A version of IMAP that is tunneled through SSL for increased security. For a description of IMAP see *Internet Access Message Protocol*.

Small Computer System Interface Commonly refered to as SCSI, is an industry standard I/O bus for high speed data transfer.

SPOP3 Is a version of the POP3 protocol that is wrapped in the SSL protocol for increased security. For a description of POP3 see *Post Office Protocol*.

SSH See *Secure Shell.*

SSL See *Secure Socket Layer.*

superuser An informal name for ROOT.

swap A swap partition is a physical hard drive partition. A Linux system utilizes swap space when system RAM starts to fill and it is necessary for more RAM. However, swap is signfigantly slower than system RAM and is not a replacement for RAM.

symlink Symbolic link. An entry in a directory that is not a file, but contains the name of another file that should normally be accessed instead. Contrasts a hard link.

trusted host A trusted host refers to a network computer or device that can be trusted. Generally these are internally controlled boxes and all boxes on the outside are untrusted.

Umask A setting in a Unix process that modifies the permissions on newly created files. It is generally represented as a three-digit octal number that will be logically ANDed against the mode 666 (rw-rw-rw). Execute bits are not on newly created files in any case.

Unix The operating system after which Linux is modeled. Although often used to refer to any operating system that provides features and programming interfaces that emulate Unix, the term is a trademark legally held by The Open Group.

user key See *host key*.

virtual interface A virtual interface is a non-existent interface that binds itself to a real interface. This virtual interface can be assigned its own IP address and will access the network through the real interface its bound to. For example interface eth0 can have eth0:X bound to it, 'X' being replaced with the virtual interface number.

virtual memory Memory beyond what is actually available, but which programs believe is actually available memory in the system. See *swap*.

Virtual Private Network Allows remote computers to connect to a common network via a medium such as the Internet as if the remote computer was locally connected to the network in a secure manner.

VPN See *Virtual Private Network.*

zone transfer A zone transfer is when a secondary name server, also sometimes
referred to as a slave server, for a zone gets the zone data from another name
server that is authoritative for the zone, called its master server. When a
secondary name server starts up, it contacts its master server and requests
a copy of the zone data for which it is responsible, storing it in the event a
request is made for information in that zone.

G REFERENCES

1. Albitz, Paul & Liu, Cricket. *DNS and BIND*, Third Edition. O'Reilly & Associates, Inc. 1998.

2. Carling, M, Degler, Stephen, and Dennis, James. *Linux System Administration*. New Riders Publishing, 2000.

3. Mark Grennan. Firewall and Proxy Server HOWTO, Feb. 26, 2000. http://www.linuxdoc.org/HOWTO/Firewall-HOWTO.html. Copyright Mark Grennan, 2000

4. Garfinkle, Simson and Spafford, Gene. Practical Unix & Internet Security, 2nd Edition.
O'Reilly & Associates, Inc. 1996

5. Hunt, Craig. *TCP/IP Network Administration*. O'Reilly & Associates, Inc. 1993

6. Laurie, Ben & Lauri, Peter, Apache The Definitive Guide, Second Edition, O'Reilly & Associates, Inc.. 1999.

7. Welsh, Matt and Kaufman, Lar, Running Linux, Second Edition, O'Reilly & Associated, Inc.. 1996

8. Dave Wreski and Kevin Fenzi, *Linux Security How-to*. http://www.linuxsecurity.com/docs/HOWTO/Security-HOWTO/, 2000

9. Wreski, Dave. *It's a Bad Bad Bad world! But Understanding the ABC's of Linux Security Can Make It Better!*. Linux Magazine, October 1999, Vol 1, Num 6, pg 31

10. Wreski, Dave. *System Security*. Linux Magazine, October 2000, Vol 2, Issue 10, pg 34.

11. Yarger, Randy Jay, Reese, George & King, Tim. MySQL & mSQL. O'Reilly & Associates, Inc. 1999

12. Zwicky, Cooper & Chapman. Building Internet Firewalls, June 2000. Copyright O'Reilly & Associates, Inc. 2000.

13. Ziegler, Robert L. *Linux Firewalls*. New Riders Publishing, 2000.

14. Zwicky, Elizabeth D., Cooper, Simon, & Chapman, D. Brent. *Building Internet Firewalls*. O'Reilly & Associates, Inc. 2000.

Index